On Sensing, Agility, and Computation Requirements for a Data-gathering Agile Robotic Vehicle

Fangchang Ma Sertac Karaman

arXiv:1704.02075v1 [cs.RO] 7 Apr 2017

***Abstract*—We consider a robotic vehicle tasked with gathering information by visiting a set of spatially-distributed data sources, the locations of which are not known *a priori*, but are discovered on the fly. We assume a first-order robot dynamics involving drift and that the locations of the data sources are Poisson-distributed. In this setting, we characterize the performance of the robot in terms of its sensing, agility, and computation capabilities. More specifically, the robot's performance is characterized in terms of its ability to sense the target locations from a distance, to maneuver quickly, and to perform computations for inference and planning. We also characterize the performance of the robot in terms of the amount and distribution of information that can be acquired at each data source. The following are among our theoretical results: the distribution of the amount of information among the target locations immensely impacts the requirements for sensing targets from a distance; performance increases with increasing maneuvering capability, but with diminishing returns; and the computation requirements increase more rapidly for planning as opposed to inference, with both increasing sensing range and maneuvering ability. We provide computational experiments to validate our theoretical results. Finally, we demonstrate that these results can be utilized in the co-design of sensing, actuation, and computation capabilities of mobile robotic systems for an information-gathering mission. Our proof techniques establish novel connections between the fundamental problems of robotic information-gathering and the last-passage percolation problem of statistical mechanics, which may be of interest on its own right.**

***Keywords*—Motion planning, stochastic environments, nonequilibrium statistical mechanics.**

I. Introduction

The Unmanned Aerial Vehicle (UAV) technology thrived during the last decade. Today, the diverse UAV market offers products with a wide range of size, power, speed, endurance, agility, and payload capacity properties, and with staggering granularity. The ability to produce and utilize UAVs with such diversity has made a number of new civilian applications commercially viable, particularly in the agriculture, security, humanitarian assistance, and disaster response domains, as evidenced by the large and increasing number of technology companies that aim to utilize UAVs in these domains.

Developing the enabling hardware, algorithms, and software still remains one of the prominent challenges. However, given the diversity of the UAV market today and the range of the potential applications, new research challenges emerge around the system design aspects, for instance, to answer the question:

The authors are with the Department of Aeronautics and Astronautics, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, 02139

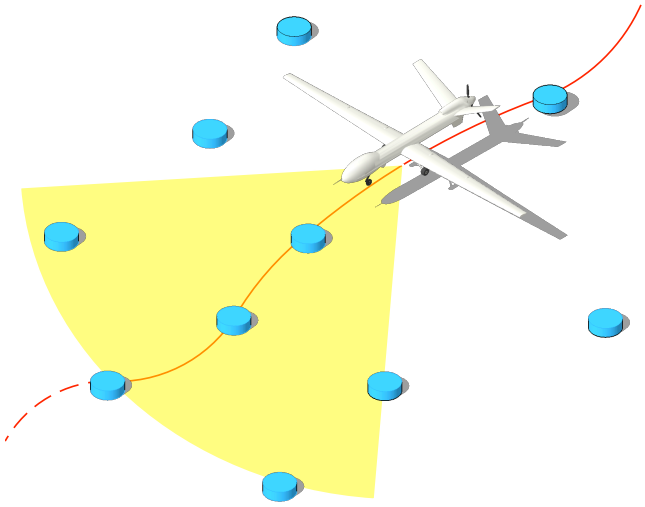

Fig. 1. An illustration of the vehicle navigating in a stochastic reward field. The blue cylinders represent the target locations. The yellow region represents the target-detection region attached to the vehicle. The locations of all targets in this range are known to the vehicle. By visiting these target locations, the vehicle can collect the reward assigned to them, as illustrated by the vehicle trajectory in red.

How can we choose the best UAV system that can address a given UAV mission, with provable guarantees on performance?

For example, consider the persistent monitoring of an agricultural field. Suppose the UAVs are tasked with locating and picturing the potential anomalies, where the locations of anomalies are not known to the UAV *a priori*, but they are discovered on the fly. Once the location with a potential anomaly is detected from a distance, *e.g.*, after that location enters the robot's vision range, the observing UAV may choose to fly over for a close-up picture to confirm the anomaly. Then, given the size of the field and how frequently the potential anomalies arise, how shall we choose the sensing, agility, and computation capabilities of the UAVs, in order to ensure that anomalies can be detected before they adversely affect the field? Note that the choice of these properties ultimately determines the the size, power, endurance, speed, and payload capacity of the UAV that shall be utilized in this mission.

Currently, such design problems are addressed in an ad-hoc manner, for instance through guesswork, or at best through extensive simulation studies. Unfortunately, rigorous mathematical tools that provide valuable insight into these system design problems still remain largely unavailable, with the exception of a few very recent results which we discuss below.

In this paper, we focus on a problem that involves informa-

tion gathering, which arises in a number of UAV applications. We consider the design of sensing, agility, and computation capabilities of a robotic vehicle that is tasked with gathering information from spatially-distributed data sources. Specifically, we prove a number of results that characterize the performance of the system in terms of *(i)* the distance from which the robot can locate the data sources, *(ii)* the agility of the robot, *(iii)* the onboard computational capabilities of the robot, and *(iv)* the amount and distribution of information at the field. These results characterize certain sensing, agility, and computation requirements for UAVs designed to execute a certain class data gathering missions. As a result, they provide insight into how to best select UAVs for the mission at hand.

Our technical approach is to reduce a certain class of data gathering problems to a *maximum reward collection problem*. First, we show that a discrete version of this maximum reward collection problem is closely related to a widely-studied model in statistical mechanics, namely the *last-passage percolation* model. Then, we analyze this problem utilizing certain results available in the statistical mechanics literature. Finally, we extend our results in the discrete domain to continuous spaces.

A. Related Work

There are a number of relevant problems studied in the literature, including the traveling salesman problem, the vehicle routing problem, the orienteering problem, and the persistent monitoring problem.

1) Traveling Salesman Problem (TSP): The TSP problem has been widely studied in the theoretical computer science and the operations research literature [1]–[4]. Given a list of n cities and the distances between every pair of them, the TSP problem is to find the shortest tour that visits each of the cities once before returning to the starting city. Although the problem is NP-hard in general, practical heuristics [5], polynomial-time approximations [6], branch-and-bound techniques [7] and learning algorithms [8] have been developed. In relation to the present paper, Beardwood *et al.* [9] consider the stochastic TSP problem where the n points are independently and uniformly distributed over a bounded region. More recently, the traveling salesperson problem for the Dubins vehicle (DTSP), *i.e.*, a non-holonomic vehicle that is constrained to move along planar paths of bounded curvature, was also studied [10], [11].

2) Vehicle Routing Problem (VRP): VRP is a generalization of the TSP problem [12]–[15]. In the VRP, a fleet of vehicles, which start from a central depot, seek to service customers. Similar to TSP, a number of exact and approximate algorithms have been proposed for solving this problem [16]–[20]. A version of the VRP problem for single vehicle is due to Wilson and Colvin [21], where customer requests appear dynamically. Dynamic and stochastic routing problems are also studied as an extension of their deterministic counterparts [22]. For example, Bertsimas and Ryzin [23] consider the the dynamic traveling repairman problem (DTRP), in which the demands for service arrive according to a Poisson process, and they provide a lower bound on the average total waiting time of demands.

The underlying maximum-reward motion problem that we study also has a stochastic and dynamic nature, since the targets are stochastically placed and they are encountered in a dynamic manner as the targets get in the sensing range of the vehicle. However, the maximum-reward motion problem differs from the stochastic and dynamic versions of the TSP and the VRP in a number of ways. Most importantly, in the maximum-reward motion problem, the vehicle is governed by dynamics with drift and it is *not* constrained to visit all of the target locations. Moreover, the dynamic nature arises from the vehicle's motion, rather than as a spatio-temporal process that is independent of the motion of the vehicle. Also, the reward placed on the targets is an essential part of the problem. Since, TSP and VRP problems constrain the vehicle to visit all of the targets, placing reward on the targets does not change the solution. Hence, we believe that the maximum-reward motion problem introduced in this paper is a fundamental stochastic dynamic decision-making problem on its own right. Its differentiating key attribute is the following: The drift affecting the vehicle does not allow visiting all of the targets, and the vehicle must choose targets based on the reward.

3) Orienteering Problem (OP): The OP originates from orienteering, an outdoor sport where a set of locations are specified. Scores are associated with these locations, and competitors seek visit a subset of these locations within a fixed amount of time in order to maximize the total score. The orienteering problem is closely tied to the traveling salesman problems (TSP) and the vehicle routing problems (VRP), and has been studied extensively in the operations research community. Its variants are also known as the maximum collection problem [24], [25], selective traveling salesman problem [26], traveling salesman problems with profits [27], and the bank robber problem [28]. This problem has been shown to be NP-hard [29], and the recent paper [30] provides a complete survey on the literature of the OP, including their variants, applications and solutions. Several heuristic methods have been proposed to tackle the OP since the early days, including deterministic and Monte-Carlo based heuristics [31], a so-called center-of-gravity heuristic [29], and many more [32], [33]. Other algorithms are also presented in the literature, including branch-and-cut algorithms [34] for finding the optimal solution, optimization-base methods [35], and search-based algorithms [36]. Researchers have also extended the OP to the case with multiple vehicles, also known as the Team Orienteering Problem [37]–[42].

The difference between the orienteering problem and this maximum-reward problem is three-fold. Firstly, in the maximum-reward problem the motion of robot is governed by its dynamics (and drift). Secondly, the OP problem is deterministic, and in comparison the maximum-reward problem has stochastic and dynamic nature arising from the robot's motion. Finally, the OP problem considers a finite-distance path (or limited-time path); whilst in this maximum-reward problem, our analysis focuses on the scaling of the performance.

4) Persistent Monitoring: The data-gathering problem which we consider is also similar to the recently-proposed persistent monitoring problems that seek to generate an optimal control strategy for a team of agents to monitor a dynamically changing environment. There exist many different formulations and applications of the persistent monitoring problem [43],

[44]. For example, a persistent monitoring task with the objective to minimize an uncertainty metric is considered in [45]–[47]. A different formulation of the persistent monitoring problem is studied in [48], [49], where robots move along fixed paths in a changing environment. In [50], [51], Alamdari *et. al* consider path planning problem for a robot to monitor a known set of features of interest, where the environment is represented as a weighted graph. In a more recent work [52], Yu *et al.* focus on a stochastic model of occurrence of events, where the precise occurrence time is not known *a priori*.

The data-gathering problem which we study differs from the problems in these references with its stochastic and dynamic nature. Moreover, our analysis fundamentally differs from these references, since we focus mainly on the asymptotic performance analysis of an optimal algorithm for a fundamental problem, while the references develop complex algorithms, often based on mathematical programming.

The particular special case that we focus on is similar to the dynamic boundary guarding problem [53], [54] studied by Smith *et. al*, where the service robot they study is constrained on a line segment while our robot can move freely along the x_2 dimension so long as the speed is bounded. Our work differs in two key aspects. First, our problem definition is more general in a number of aspects. We consider problems where the agent can move freely, whereas the aforementioned references consider a vehicle confined to a box. We attach a reward to each target and aim to maximize total reward, while the said references references maximize the number of targets visited, which is recovered when all reward is equal in our problem. Second, our proof techniques utilize techniques from statistical mechanics, specifically the last-passage percolation problem, whereas they base their analysis on the methods of combinatorial optimization in stochastic domains. In fact, the generality we provide comes with the novel connections we establish with the last-passage percolation problem.

5) Cyber-Physical Systems: Our work is in the same domain of the recent work exploring the theory of cyber-physical systems design and co-design [55]–[57]. This recent line of work aims to cast the design problem that arise in cyber-physical systems and robotics as optimization problems. Our work is complementary, and it can be considered in this emerging field of design.

It is worth noting that a dual of the maximum-reward problem has been studied [58]. The authors investigate high-speed navigation through a randomly-generated obstacle field, where only the statistics of the obstacles are given. They show that with a simple dynamical model of the vehicle, the existence of an infinite collision-free trajectory through the environment exhibits a phase transition, *i.e.*, there is an infinite collision-free trajectory almost surely when the speed is below a threshold and it will collide with some tree eventually otherwise. In [59], they show that a planning algorithm based on state lattices can navigate the robot with limited sensing range. A similar problem is also studied in [60].

6) Percolation Theory and Queues in Tandem: Finally, our work is closely tied to the literature of percolation theory (specifically, last-passage percolation [61]–[66]) and related research on queues in tandem [67], [68]. A complete survey on last-passage percolation model with general weights can be found in [69]. It can be shown that under certain technical assumptions the two models are equivalent, and results from the percolation theory have been applied in systems of queues in tandem. Recent work by Somanath *et al.* also applies similar models to control theory and robotics [70].

B. The Maximum-reward Motion Problem

As we briefly surveyed above, there are a number of problems that have been studied towards understanding missions involving autonomous vehicles in data gathering applications. However, we observe that there are no foundational problems that include sensing, agility, and computation properties at the same time in a stochastic environment. We close this gap by introducing a new foundational problem. Different from the existing literature, this problem involves a vehicle moving with drift in an environment where the tasks are distributed randomly and each target is associated with a random reward. The vehicle is endowed with a limited sensing range. In this foundational problem, the vehicle is tasked with collecting the maximum reward as it navigates through its environment. This foundational problem, which we call the *maximum-reward motion problem*, is illustrated in Figure 1. We present this problem in detail in Section II-A. We present an important special case in Section II-B, which we believe is the simplest version of the problem that still maintains its core properties. In Sections III and IV, we analyze this new foundational problem.

This foundational problem is connected with problems involving data gathering with agile robotic vehicles. We outline these connections briefly in Section II-C, and then we devote Section VI to describe this connection in detail.

C. Contributions

A preliminary version of this paper appeared in the Workshop on Algorithmic Foundations of Robotics [71], where we introduced some of the analysis for the discrete lattices presented in Section III. However, the other results in this paper, including all results in Section IV, are new. The contribution of this paper is three-fold. First, we formulate the maximum-reward motion problem, which serves as a novel mathematical foundation for the analysis of a class of data-gathering problems in robotics. Second, we provide a rigorous analysis of the robot performance, given its sensing, actuation and computation capabilities. This is achieved by establishing connections with the last-passage percolation problem in statistical mechanics. Third, we apply our results to gain insights for the design of UAV systems.

Our thorough analysis reveals a number of insights, which are explained in detail in Section VI. In particular, we find:

- The vehicle can navigate almost optimally, *i.e.*, as if it had infinite sensing range, even with very little sensing range (required sensing range scales only logarithmically with increasing mission length), when the value of information is bounded almost surely for each target location; on the contrary, when the value of information on each target is Pareto distributed (a heavy tailed distribution),

TABLE I. SUMMARY OF RESULTS

	Light-tailed distributed rewards (*e.g.*, bounded, exponential, geometric, gaussian)	Heavy-tailed distributed rewards (*e.g.*, Pareto, Cauchy, Student's t distributions)
Optimal unit-distance mean reward $R_2^*(L)$ (with unlimited sensing range)	$R_2^*(L)$ converges to a finite constant R_2^*.	$R_2^*(L)$ grows unbounded, *i.e.*, $R_2^* = \infty$. For example, for Pareto distributions with parameter α, $R_2^*(L) = O(L^{2/\alpha - 1})$
Required sensing range S for mission length L	$S = O(\log(L))$	$S = O(L)$
Impact of robot agility α on mean reward	$R_2^* = O(\alpha^{1/2})$	
Computational requirement for motion planning	$C_P = O\left(\lambda^2 \alpha^2 S^3 v\right)$	
Computational requirement for inference tasks	$C_I = O\left(\sqrt{\alpha}\, v\right)$	

the required sensing range that is almost as large as the mission length to perform optimally.

- The impact of agility on the performance of the agent can be completely characterized for the simple example that we consider. In our metrics, performance increases with increasing agility, but with diminishing returns.
- The computation requirements can also be completely quantified in terns of the sensing range and the agility of the vehicle. We find that the computational effort devoted to planning increases faster than that devoted to inference, as various parameters increase.

A summary of results can be found in Table I.

D. Organization

We formalize the maximum-reward motion problem in Section II. We introduce and analyze a discrete version of the problem in Section III. With the help of the results for the discrete case, we study the continuous problem in Section IV. In Section V, we provide the results of simulations that support our theoretical results. Finally in Section VI, we present applications of maximum-reward motions to a sensor selection problem as well as a design problem that involves a network of UAVs and unattended ground sensors.

II. PROBLEM FORMULATION

This section is devoted to a formal definition of the problem. For this purpose, we first define the problem of collecting maximum reward in a stochastic reward field in its most general form. Second, we introduce an important special case, which this paper focuses on. Finally, as an instance of this problem, we introduce an inference problem involving mobile robotic vehicles tasked with data gathering.

A. Problem 1: Maximum-reward Motion for Vehicles with Drift Operating in a Stochastic Environment

Consider a robotic vehicle navigating in a stochastic environment, where the locations of targets are distributed randomly and each target location is associated with a random reward value. The precise locations of all of the targets are unknown to the robot *a priori*. Instead, the vehicle discovers the target locations and the reward associated with the targets on the fly. To model this phenomenon, we consider a target-detection region attached to the vehicle. When the targets get inside the detection region of the robot, the locations of the targets and the reward associated with them become known to the robot. The vehicle can then choose which locations to visit and collect the reward associated with these visited targets.

Note that, when subject to differential constraints involving substantial drift, the vehicle must visit the most valuable targets in the direction of drift selectively, in order to maximize the total reward it collects. This often comes at the expense of skipping some of the target locations, for instance, those that are orthogonal to the drift direction. See Figure 1.

In this section, we present the reward collection problem in a general form. In the next section, we introduce a special case that captures all key aspects of the problem. This special case is also analytically tractable. In particular, we can derive the aforementioned fundamental limits for this special case.

The online motion planning problem is formalized as follows in its most general form:

Dynamics: Consider a mobile robotic vehicle that is governed by the following equations:

$$\begin{aligned} \dot{x}(t) &= f(x(t), u(t)), \\ y(t) &= g(x(t)) \end{aligned} \tag{1}$$

where $x(t) \in X \subset \mathbb{R}^n$ represents the state, $u(t) \in U \subset \mathbb{R}^m$ represents the control input, $y(t) \in \mathbb{R}^2$ is the position of the robot on the plane where the targets lie, X is called the state space, and U is called the control space. A state trajectory $x : [0, \tau] \to X$ is said to be a *dynamically-feasible state trajectory* and $y : [0, \tau] \to \mathbb{R}^2$ is said to be a *dynamically-feasible output trajectory*, if there exists $u : [0, \tau] \to U$ such that u, y, and x satisfy Equation (1) for all $t \in [0, \tau]$.

We are particularly interested in the case when the robot is subject to drift, for instance, when the robot can not come to a full stop instantly or can not even substantially slow down.[1] Examples include fixed-wing airplanes, racing cars, large submarines, and speed boats. In Section II-B, we will present a dynamic system model, which we believe is the simplest model that captures this drift phenomenon.

Targets and reward: The target locations and the reward associated with the targets are assumed to be generated by a stochastic marked point process.[2] A marked point process is defined as a random, countably-infinite set of pairs $\{(p_i, m_i) :$

[1] Let us note at this point that "dynamical systems with drift" can be defined precisely, for instance, through differential geometry [72]. However, we will not need such differential-geometric definitions in this paper, since we focus on a particular system with drift (introduced in Section II-B), and we leave the generalization to other drift systems to future work.

[2] Strictly speaking, stochastic point processes are formalized using counting measures [73]. For the sake of the simplicity of the presentation, we will avoid these measure-theoretic constructs, and instead we will use the simpler notation adopted by Stoyan et al. [74].

$i \in \mathbb{N}\}$, where $p_i \in \mathbb{R}^2$ is the location of point i in the infinite plane and $m_i \in M$ is the mark associated with point i. We denote this random set by Ψ. With a slight abuse of notation, we denote by $\Psi(A)$ the number of points in a subset $A \in \mathbb{R}^2$ of the infinite plane. Given a point p of the point process, we denote its mark by $r(p)$. In our case, the locations $\{p_i\}$ of the points represent the locations of the targets, and the marks $\{m_i\}$ represent the reward associated with the targets. Hence, the mark set is the set of all non-negative real numbers, *i.e.*, $M = \mathbb{R}_{\geq 0}$. Following Stoyan et al. [74], we make the following technical assumptions: *(i)* any bounded subset of the plane contains finitely many points, *i.e.*, $|\Psi(A)| < \infty$ for all bounded measurable $A \subset \mathbb{R}^2$; *(ii)* no two points are located at the same location, *i.e.*, $p_i \neq p_j$ for all $i \neq j$, almost surely.

Target-sensing region: The locations of the targets and the reward associated with them is not known *a priori*, but is revealed to the robot in an online manner. This aspect of the problem is formalized as follows. Let $\mathcal{P}_\Psi(\cdot)$ denote the target-detection region of the robot that associates each state $z \in X$ of the robot with a region $\mathcal{P}_\Psi(z) \subset \mathbb{R}^2$. When the robot is in state $z \in X$, it is able to observe only those targets that lie in the set $\mathcal{P}_\Psi(z)$. That is, $\{(p_i, m_i) \in \Psi : p_i \in \mathcal{P}_\Psi(z)\}$ is the set revealed to the robot when it is in state z.

Task: The robot is assigned the task of collecting maximum total reward, subject to all of the constraints outlined above. We formalize this objective of the problem as follows. Suppose the stochastic marked point process that represents the targets is defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the sample space, $\mathcal{F}$ is the σ-algebra, and $\mathbb{P}$ is the probability measure. Define $\mathcal{F}_t$ as the σ-algebra generated by the random variables $\cup_{t' \in [0,t]} \mathcal{P}_\Psi(x(t'))$. A *feasible control policy* is a stochastic process $\mu = \{u(t) : t \in [0, \tau]\}$, such that $u(t)$ is defined on $(\Omega, \mathcal{F}_t, \mathbb{P})$, for all $t \geq 0$.[3] This definition implies that the control policy depends only on the locations and the reward of the targets that are detected by the robot up until time t, but not at times greater than t. Note that the control policy may depend on the statistics of the stochastic marked point process, if any statistics are known *a priori*.

Given a control policy $\mu = \{u(t) : t \in [0, \tau]\}$, let us denote the collection of resulting state trajectories by $\{x_\mu(t) : t \in [0, \tau]\}$ and the collection of the output trajectories by $\{y_\mu(t) : t \in [0, \tau]\}$, which are stochastic processes defined on the same probability space as the control policies.

Let $Y(\mu; \psi)$ denote the set of targets visited by the robot under control policy μ and when the realization of the stochastic marked point process for the targets is ψ, *i.e.*,

$$Y(\mu; \psi) = \big\{p \in \psi : y_\mu(t) = p \text{ for some } t \in [0, \tau]\big\}.$$

With a slight abuse of notation, let $T(\mu; \psi)$ denote the total reward collected by the robot when it visits these targets, *i.e.*,

$$T(\mu; \psi) = \sum_{p \in Y(\mu; \psi)} r(p).$$

[3]We omit some of the measure-theoretic details when defining the control policy. Our definition matches the definition of an adapted control policy introduced by Kushner [75].

Then, the *maximum-reward motion problem* is to find a control policy μ such that the total reward $T(\mu; \psi)$ is maximized for all realizations of ψ of the point process Ψ.

We stress that an algorithmic solution to this problem (*i.e.*, computing such a policy) is often simple, particularly when the point process Ψ is completely random (*e.g.*, a Poisson process). Instead of designing new algorithms, we are more interested in analytically deriving the maximum reward that can be achieved by an optimal algorithm. Such analyses may allow robotics engineers to design robotic systems, *e.g.*, by choosing the sensing, actuation, and computation capabilities of the robots, such that they best fit the application at hand.

B. Problem 2: Maximum-reward Motion with Lipschitz-continuous Paths in a Poisson Target Field with I.I.D. Rewards

In this section, we present a two-dimensional special case. We believe this is the simplest case that captures all key aspects of the problem we presented above, namely: *(i)* the dynamics with drift, *(ii)* the stochastic nature of the environment, and *(iii)* the online (dynamic) nature of the task. This simple case is particularly relevant to the motivational example presented in Section I on information gathering. Furthermore, this case is also analytically tractable.

Dynamics: Let $X = \mathbb{R}^2$, and define the dynamics governing the robot with the following ordinary differential equation:

$$\begin{aligned} \dot{x}_1(t) &= v, \\ \dot{x}_2(t) &= u(t), \end{aligned} \tag{2}$$

where $[x_1(t)\ x_2(t)] \in \mathbb{R}^2$ denotes the state of the robot, v is a constant, and $|u(t)| \leq w$ is the control input. This robot travels with constant speed v along the longitudinal direction (x-axis) and with bounded speed u in the lateral direction (y-axis). We define the *agility* of the robot as $\alpha = w/v$. The larger this number α, the more maneuverable the robot is.

Targets and reward: The target locations are generated by a two dimensional *Poisson point process* with intensity λ. That is, the number of targets $\Psi(A)$ for any region $A \in \mathbb{R}^2$ follows a Poisson distribution, *i.e.*, $\Psi(A) \sim Poi(\lambda\,|A|)$. The reward associated with each target is chosen from a common distribution independently. Let $r(p)$ denote the reward associated with the target at location $p \in \mathbb{R}^2$. Then, $\{r(p_i) : i \in \mathbb{N}\}$ are independent identically distributed random variables.

Target-sensing region: The robot has a fixed target-sensing range m. That is, when the robot is at state $x(t) = [x_1(t)\ x_2(t)]$, it obtains the target's information, namely its location $p = [p_1\ p_2]$ and the associated reward m, for all targets located in

$$\mathcal{P}_\Psi(x(t)) = \big\{p \in \mathbb{R}^2 \,:\, 0 \leq p_1 - x_1(t) \leq m \text{ and } (p, m) \in \Psi\big\}.$$

C. Data Gathering and Inference with Gaussian Noise

We established the maximum-reward motion in stochastic environments as a foundational problem. How is this problem related to data gathering with an agile robotic vehicle?

In this section, we relate the maximum-reward motion problem with a special data-gathering problems as follows. We consider a robot tasked with estimating a fixed scalar θ

from spatially-distributed measurements corrupted by Gaussian noise. As the robot navigates through the two-dimensional environment, it observes various candidate locations where measurements can be taken. The robot also observes the "value of the information" that these locations may potentially provide, if the robot visits that location and measures θ. Based on the locations and their potential value of information, the robot must decide a subset of the locations to visit, and collect noisy measurements of the variable θ at each of those locations.

The higher the value of information, the less noisy will be the observation at that location. Let the prior belief on θ be Gaussian-distributed with mean μ_0 and variance $\frac{1}{\beta_0}$, *i.e.*,

$$\theta \sim \mathcal{N}(\mu_0, 1/\beta_0).$$

Let the likelihood function of measurement y_i of θ also be Gaussian distributed, centered at θ with variance $\frac{1}{\beta_i}$, *i.e.*,

$$y_i|\theta \sim \mathcal{N}(\theta, \frac{1}{\beta_i}). \tag{3}$$

Notice that β_i is the precision of measurement y_i. Given sensor measurements $\mathbf{y} = [y_1, y_2, \ldots, y_n]$, we can derive the posterior probability of θ conditioning on $\mathbf{y}$ using Bayes' rule,

$$\theta|\mathbf{y} \sim \mathcal{N}(\mu_n, 1/\beta'_n),$$

with the updated mean μ_n and variance $\frac{1}{\beta'_n}$ satisfying

$$\beta'_n = \beta_0 + \beta_1 + \cdots + \beta_n. \tag{4}$$

Suppose the potential sensing locations are randomly distributed in the environment, and the robot is tasked with estimating θ, subject to the differential constraints given by Equation (2). The robot does not know the precise locations where measurements can be taken, but instead these target locations are discovered on the fly. Once a target location p_i enters the target-detection region of the robot, the robot observes the precision β_i of the corresponding measurement at p_i. If the robot chooses the visit y_i, then it will measure θ, where the measurement is corrupted with Gaussian noise of variance $1/\beta_i$.

The robot is assigned the task of estimating θ as best as possible as its navigates through the field, and its performance is measured by the variance of the posterior distribution (the lower the better). This problem is a specific instance of the maximum-reward motion problem defined in Section II-B. In this setting, the "reward" associated with sensing location y_i is precisely β_i. This problem is also motivated by the selection of unattended ground sensors, which we discuss as an application in Section VI-B. However, in Sections III and IV, where we present our main results, we will focus on the maximum-reward problem, since we believe that the maximum-reward problem represents a more general setting.

III. Preliminaries: Last Passage Percolation and the Analysis of Motion on State Lattices

In this section, we develop some preliminaries. Specifically, we introduce a discrete problem that approximates the continuous problem of Section II, and we devote this section to the analysis of this discrete problem. Many of our main results presented in the next section are obtained as the limiting cases of our results for the discrete problem in this section.

The approximate problem is constructed by using a lattice-based discretization. Note that lattice-based motion planning algorithms have long been widely adopted in robotics applications [76]–[79]. These algorithms form a directed lattice in the state space of the robot and select the optimal path through this lattice. This task is often computationally efficient, making it a practical approach even for challenging problem instances.

We analyze this discrete problem and the resulting lattice-based planning algorithm, by establishing connections between this class of problems and a class of problems in non-equilibrium statistical mechanics. Roughly speaking, we view the robot as a particle traveling in a stochastic field. This perspective allows us to directly apply some of the recent results from the last-passage percolation problem [61], [64], [66], [69].

In what follows, we describe the lattice-based discretization in Section III-B. We introduce a lattice-based planning algorithm in Section III-C. We analyze the fundamental limits of the problem in Section III-D and the performance of the iterative planning algorithm in Section III-E.

A. The Last Passage Percolation Problem

A graph (V, E) is called a *d-dimensional regular lattice*, if $V = \mathbb{N}^d$, and $(v, v') \in E$ if and only if $v = (v^1, v^2, \ldots, v^d)$ and $v' = (v^1, v^2, \ldots, v^{k-1}, v^k + 1, v^{k+1}, \ldots, v^d)$ for some k. The two-dimensional regular lattice is illustrated in Figure 2a.

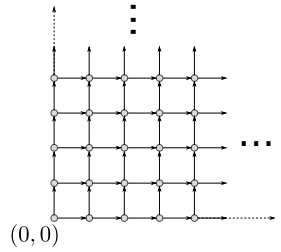

(a) The two-dimensional lattice.

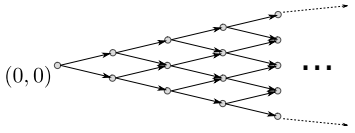

(b) Two-dimensional lattice as the discretization of Equation (2).

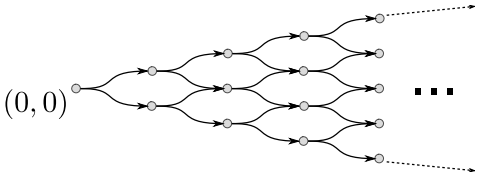

(c) A lattice with Dubins paths.

Fig. 2. The two-dimensional directed regular lattice, $\mathbb{N}^2$, is illustrated in Figure (a). An example state-lattice for a curvature-constrained Dubins vehicle, is shown in Figure (b). The latter lattice can be embedded in $\mathbb{N}^2$.

Let $|v|$ denote the distance of the vertex v from the origin, measured by the number of vertices that any path between these two vertices has to visit. In other words, for the vertex

v as a d-dimensional vector of natural numbers, *i.e.*, $v = (v_1, v_2, \ldots, v_v) \in \mathbb{N}^d$, we know that $|v| := \|v\|_1 = \sum_{j=1}^{d} v_j$.

Let $\Pi(v_0, v)$ denote the set of all possible paths that start from vertex $v_0 \in V$ and end at vertex $v \in V$. With a slight abuse of notation, let $\Pi(v_0, n)$ denote the set of all paths that start at vertex v_0 and cross exactly n vertices. When v_0 is the origin, *i.e.*, $v_0 = \mathbf{0}$, we drop it from the notation, and we simply write $\Pi(v)$ for the set of all paths that start from the origin and end at vertex v. Then, the maximum total reward starting from v_1 and reaching v on a d-dimensional regular lattice is defined as:

$$T_d^*(v_0, v) := \max_{\pi \in \Pi(v_0, v)} \sum_{v \in \pi} r(v),$$

where we write $v \in \pi$, when a path π crosses a vertex v. We denote the maximum total reward by a path that starts from v_1 and crosses at most n vertices as:

$$T_d^*(v_0, n) := \max_{\pi \in \Pi(v_0, n)} \sum_{v \in \pi} r(v).$$

Finally, when $v_0 = \mathbf{0}$ (i.e., when the paths start at the origin), we simply drop it from the notation: We write $T_d^*(v)$ and $T_d^*(n)$ for $T_d^*(\mathbf{0}, v)$ and $T_d^*(\mathbf{0}, n)$, respectively.

If a vertex v is exactly n steps away from the origin, *i.e.*, $|v_{\text{dest}}| = n$, then $\Pi(v_{\text{dest}}) \subset \Pi(n)$. Since $\Pi(v_{\text{dest}})$ is a subset, it is obvious that its maximum reward cannot exceed that of $\Pi(n)$. In other words, $T_d^*(v_{\text{dest}}) \leq T_d^*(n)$.

B. Problem 3: Maximum-reward Motion on State Lattices

In this section, we formulate a discrete problem that resembles Problem 2 of Section II-B. This new problem indeed captures a discretized version of Problem 2. However, we stress that the problem presented in this section is *not* a special case.

Dynamics: A d-dimensional directed regular state lattice $L_d = (V, E)$ is a graph that satisfies: *(i)* V is a countable set of vertices such that each vertex $v \in V$ is a state of the vehicle; *(ii)* $E \subset V \times V$ is a set edges, such that for all $(v, v') \in E$, there exists a dynamically-feasible trajectory $x_e : [0, \tau_e] \to X$ that connects v and v', *i.e.*, $x(0) = v_1$ and $x(\tau_e) = v_2$; *(iii)* L_d is isomorphic to a d-dimensional regular lattice.

Therefore, the dynamics of Equation (2) can be discretized as a two-dimensional directed regular state-lattice, as given in Figure 2b. An example state-lattice for a non-holonomic vehicle is shown in Figure 2c. Both examples are isomorphic to the two-dimensional regular lattice in Figure 2a.

Targets and reward: Each vertex $v \in V$ is associated with an independent, identically distributed random reward $r(v)$.

Target-detection range: The target-detection range is a positive number m. Any vertex reachable with a path of length m is considered to be within the target-detection range. That is, the vehicle can observe the reward $r(v)$ for all vertices that are within a distance of m to the vehicle.

Task: The vehicle is again tasked with moving through its environment and maximizing the total reward collected.

From Figure 2b, we see that Problem 3 of this section is a discrete version of the Problem 2 of Section II-B. We stress that Problem 3 of this section can be more general. For example, the state lattice in 2c represents the Dubins vehicle dynamics.

In the rest of Section III, we analyze Problem 3. This analysis will be used for deriving our main results for Problem 2. These results will be presented in Section IV.

C. Iterative Motion Planning Algorithm for State Lattices

A path on $G = (V, E)$ is a sequence of vertices, $(v_1, v_2, \ldots, v_k)$, such that consecutive vertices are connected with an edge, *i.e.*, $(v_i, v_{i+1}) \in E$ for all $i \in \{1, 2, \ldots, k-1\}$. The set of all paths on G that starts at vertex v is denoted by $\texttt{Paths}(v)$.

We consider the following motion planning algorithm. Suppose the vehicle starts at an initial state $z_{\text{init}} \in V$. In each iteration, the maximum-reward path, say $(v_1, v_2, \ldots, v_k)$, within the "visible" portion of the lattice is computed, and the vehicle executes the resulting dynamically-feasible trajectory until its end. The same procedure is repeated, after the vehicle reaches the final state $v' = x_{e_k}(\tau_{e_k})$.

We call this algorithm the *iterative lattice-based online motion planning algorithm*, which we formalize in Algorithm 1. The $\texttt{GetCurrentState}()$ procedure returns the current state of the robot, and $\texttt{Execute}(x)$ refers to the command that makes the robot follow the trajectory x. The algorithm first retrieves the robot's current state (Line 2). Subsequently, it observes the reward associated with the vertices for the visible region of the lattice (Line 3). It then searches for the optimal path over this region (Line 4). Finally, it executes this path until the vehicle reaches the end of the path (Line 5). This procedure continues for N iterations (Lines 1-5).

Algorithm 1 Iterative lattice-based online motion planning

1: **for** $t = 1, \ldots, N$ **do**
2: state $\leftarrow \texttt{GetCurrentState}()$
3: $\texttt{PerceiveEnvironment}()$
4: $\pi \leftarrow \arg\max_\pi \{T(\pi) : \pi \in \texttt{Paths}(\text{state})\}$
5: $\texttt{Execute}(\pi)$

In Line 4, the algorithm computes the maximum-weight path on a finite weighted graph. Let us note that this problem is NP-hard in general [80]. However, the problem can be solved efficiently on acyclic graphs [80].[4]

In what follows, we analyze the performance of Algorithm 1 on the 2-dimensional directed regular state-lattice (and more generally, d-dimensional regular lattices).

D. Mean Reward with Unlimited Sensing Range

In this section, we analyze the maximum reward that robot can collect, if it had infinite sensing range. The results of this section can be regarded as fundamental limits: Even when the robot has infinite sensing range, the reward it can collect is bounded by what we report in this section.

[4]Acyclic graphs arise in lattice-based motion planning, for instance, when the robot does not return to previously visited locations, *i.e.*, the robot constantly explores new regions in the environment. When lattice-based motion planning algorithms are applied to robots subject to substantial drift, the resulting lattice is also often acyclic. It is easy to show that any state lattice for the system presented in Section II-B is acyclic.

For our analysis, we focus on the maximum reward collected per unit distance traveled, which we call the *mean reward*. We define the *maximum mean reward* as follows:

$$R_d^*(n) := \frac{T_d^*(n)}{n}.$$

First, we show that the limit $\mathbf{R}_d^* := \lim_{n\to\infty} \mathbb{E}[R_d^*(n)]$, which we call *expected maximum mean reward*, is well defined.

Theorem 1. *The following holds:*

$$\lim_{n\to\infty} \mathbb{E}[R_d^*(n)] = \lim_{n\to\infty} \frac{\mathbb{E}[T_d^*(n)]}{n} = \sup_{n\in\mathbb{N}} \frac{\mathbb{E}[T_d^*(n)]}{n} \in \mathbb{R}\cup\{\infty\}.$$

The proof can be found in Appendix A. For general distributions, the value of $\mathbf{R}_d^*$ can not be computed directly. However, we can compute asymptotics for special cases of light-tailed (exponential and geometric distributions, specifically) and heavy-tailed distributions.

A distribution is light tailed if its tail is bounded by an exponentially decreasing function. For example, Gaussian, geometric, exponential and all bounded distributions are light-tailed, while Pareto, logarithmic normal, Cauchy and Student's t distributions are heavy-tailed. More precisely:

Definition 1 (See [81])**.** *The distribution F is said to be light-tailed, if it has an exponentially bounded tail,* i.e.*, for some $a, b > 0$, we have $1 - F(x) \le ae^{-bx}$, for all $x > 0$. A heavy-tailed distribution is one that is not light-tailed.*

1) Light-tailed reward distributions: The following theorem shows that, when the reward distribution is a light-tailed distribution, the maximum mean reward is bounded.

Theorem 2 (See Theorem 4.1 in [69])**.** *Suppose the rewards r are independent and identically distributed with distribution F, such that*

$$\int_0^\infty (1 - F(s))^{1/d} ds < \infty, \tag{5}$$

then $\mathbf{R}_d^*$ is finite for any $d \ge 2$, i.e.*,*

$$\mathbf{R}_d^* = \lim_{n\to\infty} \mathbb{E}[R_d^*(n)] = \lim_{n\to\infty} \mathbb{E}\left[\frac{T^*(n)}{n}\right] = c_1$$

for some finite constant $c_1 \in \mathbb{R}$.

The proof can be found in Appendix B. Note that Equation 5 is only very slightly stronger than the existence of a finite d^{th} moment, *i.e.*, $\mathbb{E}[r^d] < \infty$. Moreover, when the dimension is $d = 2$ and if rewards follow exponential or geometric distributions (both satisfy Equation 5), then $\mathbf{R}_2^*$ can be computed explicitly.

Proposition 1. *For exponential and geometric reward distributions with mean μ and variance σ^2 on a two-dimensional lattice, the expected maximum mean reward $\mathbf{R}_2^*$ can be computed explicitly*

$$\mathbf{R}_2^* = \mu + \sigma.$$

The proof is given in Appendix C.

2) Heavy-tailed reward distributions: We just showed that, for all light tailed distributions, $\mathbf{R}_2^*$ is finite. In contrast, the following theorem states that, when the reward distribution is heavy-tailed, $\mathbf{R}_2^*$ is infinite.

Theorem 3 (See Proposition 2 in [82])**.** *Suppose the rewards are independent and identically distributed and $\mathbb{E}[r^d] = \infty$, then $\mathbf{R}_d^* = \infty$.*

Here we consider a specific instance of the heavy-tailed distribution family, the *Pareto distribution*, which is commonly used to describe the allocation of wealth among individuals or distribution of income. More specifically,

Definition 2 (Pareto Distribution)**.** *The Pareto distribution with index parameters x_m and α is defined as follows:*

$$\mathbb{P}(X \le x) = \begin{cases} 1 - \left(\frac{x_m}{x}\right)^\alpha, & x \ge x_m \\ 0, & x < x_m \end{cases}$$

For a Pareto distribution, more accurate results regarding the growth rate of $T_2^*(n)$ can be obtained as follows:

Proposition 2. *Suppose the rewards r are independent and Pareto distributed with parameter $\alpha \in (0, 2)$. Then, the optimal mean reward $\mathbf{R}_2^*$ is infinite. Moreover, the growth rate of $T_2^*(n)$ is at the order of $n^{(2/\alpha)}$,* i.e.*,*

$$T_2^*(n) = O\left(n^{2/\alpha}\right),$$
$$R_2^*(n) = O\left(n^{2/\alpha - 1}\right).$$

The proof is given in Appendix D.

E. Mean Reward with Limited Sensing Range

In this section, we consider robots with a limited sensing range. Suppose the sensing range is m. To travel a distance of n vertices, we follow the best path for m steps, and then repeat this procedure, until the nth vertex is reached. We are particularly interested in comparing the reward collected with limited sensing range in this way to the reward collected with unlimited sensing range as described in the previous section.

Let T_1 denote the maximum total reward collected by a path that starts from the origin vertex, $\mathbf{0}$, and has length m, *i.e.*, $T_1 := T_d^*(\mathbf{0}, m)$. Let v_1 denote the vertex where the maximum-reward path achieving reward T_1 ends. More generally, define $T_k := T_d^*(v_{k-1}, m)$, where v_k is the vertex the path achieving reward T_{k-1} ends. Assume n is a multiple of m. Define:

$$T_{\text{iterative}}(n; m) := \sum_{i=1}^{n/m} T_i. \tag{6}$$

Finally, define the mean reward collected by the robot in n steps with sensing range m as:

$$R_{\text{iterative}}(n; m) = \frac{T_{\text{iterative}}(n; m)}{n}.$$

In our analysis, we compare mean reward $R_{\text{iterative}}(n; m)$ with $\mathbf{R}_d^*$. Recall that the former is the mean reward that the

robot can collect with limited sensing range m, and the latter is the mean reward with unlimited sensing distance.

Surprisingly, the difference in performance between the unlimited versus limited sensing range turns out to be drastically different, when the distribution of the reward is light tailed versus heavy tailed. We analyze both cases in this order.

1) Light-tailed reward distributions: Theorem 4 shows that when the reward is in the light-tailed family, iterative motion planning algorithms achieve near-optimal performance even with very limited sensing range.

Theorem 4. *Suppose the rewards r are independent, identically distributed and satisfy Equation 5. Then, for any $\delta > 0$, there exists a constant c such that $R_{\text{iterative}}(n, c\log n)$ converges to $\mathbf{R}_d^*$ in probability,* i.e.*,*

$$\lim_{n\to\infty} \mathbb{P}\left(\left| R_{\text{iterative}}(n, c\log n) - \mathbf{R}_d^* \right| \geq \delta \right) = 0.$$

The proof of Theorem 4 is provided in Appendix E. Roughly speaking, Theorem 4 implies that the robot can navigate to any vertex that is n steps away almost optimally (as if it had infinite sensing distance), even when its sensing range is only $c\log n$. This result is remarkable, as $\log n$ is much smaller than n. Our simulation results provided in Section V support our conjecture when $d = 2$.

2) Heavy-tailed reward distributions: We consider the case where the rewards follow a Pareto distribution on a two-dimensional regular lattice, *i.e.*, $d = 2$. We show that the iterative motion planning algorithm can *not* achieve a near-optimal performance with limited sensing distance $o(n)$.

Theorem 5. *Suppose the assumptions of Proposition 2 hold. Then, there exists a probability space $(\Omega, \mathcal{F}, P)$ such that when $M(n)$ is a sub-linear function of n,* i.e.*, $\lim_{n\to\infty} M(n)/n = 0$, we have*

$$\lim_{n\to\infty} \frac{R_{\text{iterative}}\left(n; M(n)\right)}{R_2^*(n)} = 0.$$

See Appendix F for the proof. Roughly speaking, Theorem 5 states that, if the sensing range of the robot is slightly less than n, then the reward collected will be much lower.

The findings of this section are summarized below.

Remark 1. *According to Theorem 4, when the reward follows a light-tailed distribution, $c\log n$ sensing range is adequate to navigate optimally (as if the robot had infinite sensing range). However, according to Theorem 5, when the reward distribution follows the Pareto law (a heavy-tailed distribution), any non-negligible limitation in sensing range leads to substantial losses in performance.*

IV. Analysis of the Continuous Problem

In this section, we return to Problem 2 (see Section II-B), namely the maximum-reward motion problem in $\mathbb{R}^2$. We study this problem, assuming (i) unit agility, (ii) infinite sensing range, and (iii) infinite computation capability. We find the necessary requirements on agility, sensing range, and computation capabilities that will allow the robot to perform close to these fundamental limits

Our analysis is based on our results presented in Section III. Specifically, we show that the continuous problem is the limiting case of the discrete Problem 3 (see Section III-B).

This section is organized as follows. In Section IV-A we analyze the fundamental limits of robot performance, given infinite sensing range with unit agility. We introduce the iterative motion planning algorithm in Section IV-B and study its performance under limited sensing range in Section IV-C for different types of reward distributions. In Section IV-D we move on to requirements on robot agility. In Section IV-E we study the computational workload for motion planning and inference tasks, respectively.

A. Fundamental Limits: The Analysis of Mean Reward with Infinite Sensing Range and Unit Agility

Let $\Pi(L)$ be the set of all feasible paths that start from the origin and travels a distance of L in the longitudinal direction, *i.e.*, the x_1 axis. Recall the assumption that the reward locations $\{p_i\}$ are generated by a Poisson point process with intensity λ. The amount of reward at each target is an i.i.d. random variable $r(p_i)$ that follows some common reward distribution.

Let $T_d^*(L)$ denote the optimal total reward collected by following any path in $\Pi(L)$ with infinite sensing distance, *i.e.*,

$$T_d^*(L) := \max_{\pi\in\Pi(L)} \sum_{p_i\in\pi} r(p_i).$$

Let $R_d^*(L)$ denote the optimal mean reward collected in the same manner, *i.e.*,

$$R_d^*(L) := \frac{T_d^*(L)}{L}.$$

Throughout this section, we assume unit robot agility $\alpha = 1$, and we analyze the total reward that the robot can collect.

The first two results for the continuous problem are extensions of Theorem 1-2.

Theorem 6 (Well-posedness of the Mean Reward)**.** *Suppose the reward locations are generated by a Poisson point process with intensity λ on $\mathbb{R}^2$. The reward associated with each target is chosen from a common distribution F independently. The robot dynamics satisfies the following ordinary differential equation:*

$$\dot{x}_1(t) = v, \quad \dot{x}_2(t) = u(t), \quad |u(t)| \leq v.$$

Then,

$$\lim_{L\to\infty} \frac{\mathbb{E}T_d^*(L)}{L} = \sup_L \frac{\mathbb{E}T_d^*(L)}{L}.$$

The proof for both Theorem 6 is given in Appendix G. For simplicity of notation we will define this optimal mean reward as

$$\mathbf{R}_2^* := \sup_L \frac{\mathbb{E}T_d^*(L)}{L}.$$

Next, we compute asymptotics for the optimal mean reward. As in the discrete case, we consider the light-tailed and heavy-tailed reward distributions separately.

1) Light-tailed reward distributions: The following theorem is the continuous counterpart of Theorem 2, and it shows that, when the reward distribution is light-tailed, the maximum mean reward is bounded.

Theorem 7 (Mean Reward Asympototics for Light-tailed Rewards)**.** *Suppose the conditions of Theorem 6 hold, and the rewards $r(v)$ are independent, identically distributed, and satisfy Equation 5. Then, $\mathbf{R}_d^*$ is finite for any $d \geq 2$,* i.e.*,*

$$\mathbf{R}^* = \lim_{n\to\infty} \mathbb{E}[R_d^*(n)] = \lim_{n\to\infty} \frac{\mathbb{E}[T^*(n)]}{n} = c_2$$

for some finite constant $c_2 \in \mathbb{R}$.

Interested reader please refer to the proof for Theorem 1.2.1 in [83], which applies the subadditive ergodic theorem.

2) Heavy-tailed reward distributions: In contrast, the following theorem shows that, when the reward distribution is heavy-tailed, $\mathbf{R}_2^*$ is infinite. This theorem is the continuous counterpart of Theorem 3.

Theorem 8 (Mean Reward for Heavy-Tailed Rewards)**.** *Suppose the conditions of Theorem 6 hold and the rewards satisfy $\mathbb{E}[r^2] = \infty$. Then, the optimal mean reward $\mathbf{R}_2^*$ is infinite.*

The proof is in Appendix H. Similar to the case with discrete lattices, more accurate results can be derived for the Pareto distributions.

Proposition 3 (Mean Reward Asympototics for Pareto Rewards)**.** *Suppose the conditions of Theorem 6 hold, and the rewards r are Pareto-distributed with parameter $\alpha \in (0, 2)$. Then, the optimal mean reward $\mathbf{R}^*$ is infinite. Moreover, the growth rate of $T_2^*(L)$ is order $L^{(2/\alpha)}$,* i.e.*,*

$$T_2^*(L) = O\left(L^{2/\alpha}\right),$$
$$R_2^*(L) = O\left(L^{2/\alpha - 1}\right).$$

The proof is given in Appendix I.

B. The Iterative Motion Planning Algorithm

Similar to the planning algorithm on discrete lattices, in the continuous space the planning algorithm proceeds in an iterative manner. Suppose the robot starts at an initial state z_{init}. First, the best feasible trajectory $x_e : [0, T_e] \to X$ within the "visible" region of the lattice is computed, and the robot follows this dynamically-feasible trajectory until the end. After the robot completes this trajectory, the same procedure is repeated. This algorithm is formalized in Algorithm 2.

Let `PerceiveEnvironment()` (Line 4) be a procedure that returns the set of targets/rewards that are visible to the robot. The robot then computes the optimal path within the set of trajectories `Paths`(state) to maximize the total reward collected (Line 5). In this problem, `Paths` $= \{\pi : \dot{x_1} = v, |\dot{x_2}| \leq w\}$. The procedure `Execute`($\pi$) (Line 6) commands the robot to move along the planned path $\pi : [0, m/v] \to X$ and returns the total reward collected along this path. After completion of this command, the entire procedure is repeated until time $distance_x$ is greater than travel distance L (Lines 2-8).

Algorithm 2 Receding-horizon online motion planning

1: `distance`$_x \leftarrow 0$
2: **while** `distance`$_x < L$ **do**
3: state $\leftarrow$ `GetCurrentState()`
4: `PerceiveEnvironment()`
5: $\pi \leftarrow \arg\max\{T(\pi) : \pi \in$ `Paths`(state)$\}$
6: $T_i \leftarrow$ `Execute`(π)
7: $Q \leftarrow Q + T_i$
8: `distance`$_x \leftarrow$ `distance`$_x + m$

C. Sensing Requirements

We denote the sensing distance by S. Define T_i is the amount of total reward collected during the i^{th} iteration of Algorithm 2, and let $T_{\text{iterative}}(L; m)$ denote the total reward collected with Algorithm 2 throughout the entire mission, *i.e.*,

$$T_{\text{iterative}}(L; S) := \sum_{i=1}^{L/S} T_i.$$

Define the mean reward collected by the Algorithm 2 by

$$R_{\text{iterative}}(L; S) := \frac{T_{\text{iterative}}(L; S)}{L}.$$

In this section, we analyze the mean reward collected by the algorithm for two different reward distributions: (i) when the rewards are almost-surely bounded, (ii) when the rewards follow the Pareto distribution. The former is a light-tailed distribution, whereas the latter is a heavy-tailed distribution.

1) Light-tailed reward distributions: The following result extends Theorem 4 and shows that the receding horizon algorithm still has near-optimal performance even in the continuous problem, when the sensing distance m is at the order of $\log L$.

Theorem 9 (Sensing range requirements for light-tailed rewards)**.** *Suppose the conditions of Theorem 6 hold. Then, for any $\delta > 0$, there exists some constant $c_3 > 0$ such that*

$$\lim_{m\to\infty} \mathbb{P}\left(\, |R_{\text{iterative}}(L; c_3 \log L) - \mathbf{R}_2^*| \;\geq\; \delta \,\right) \;=\; 0.$$

2) Heavy-tailed reward distributions: The following theorem shows that, when the rewards follow the Pareto distribution (a heavy-tailed distribution), then the

Theorem 10 (Sensing range requirements for Pareto rewards)**.** *Suppose the assumptions of Proposition 3 hold. Then, there exists a probability space $(\Omega, \mathcal{F}, P)$ such that, for any sublinear function $S(L)$,* i.e.*, $\lim_{L\to\infty} S(L)/L = 0$, we have*

$$\lim_{L\to\infty} \frac{\mathbb{E}[R_{\text{iterative}}(L; S(L))] - \mathbb{E}[R_d^*(L)]}{L^{(2/\alpha)-1}} = c_4$$

for some positive constant c_4.

To avoid repetition we don't provide a detailed proof for Theorem 9-10, since the proof techniques are identical to Theorem 4-5.

D. Agility Requirements

In this section, we examine how agility impacts the performance of the robot, measured by the total reward collected. The main result of this section is the following:

Theorem 11 (Agility Requirements)**.** *Suppose the reward locations are generated by a Poisson point process with intensity λ on $\mathbb{R}^d$. The robot dynamics satisfies the following ordinary differential equation:*

$$\dot{x}_1(t) = v, \quad \dot{x}_2(t) = u(t),$$

where $|u(t)| \leq w$. Then for any finite $L > 0$, there exists a constant $c_5 > 0$ that is independent of A (but depends on L) such that

$$\mathbb{E}[R_d^*(L)] = c_5\sqrt{\alpha},$$

where $\alpha = w/v$ is the agility of the robot.

E. Computation Requirements

In this section, we analyze the computational capabilities required on board the robot to perform the inference and planning tasks for data gathering. Specifically, we first analyze the amount of computational operations required to run the planning algorithm presented in Algorithm 2. We then analyze the amount of computational operations required to run inference algorithms, *e.g.*, in the setting described in Section II-C.

a) On motion planning: The planning algorithm, presented in Algorithm 2, is called periodically. Recall that S is the sensing range of the robot and L is the length of the mission. Then, the planning algorithm is called exactly L/S times over the course of the mission. The planning procedure itself is a dynamic programming algorithm that computes the optimal path on an acyclic graph with N nodes, where N is the number of targets that are within the sensing range of the robot when the algorithm is called. The dynamic programming requires $O(N^2)$ steps. The expected value of the number can be computed as follows. The area of the target-detection region is of order αS^2. Since the target locations are Poisson distributed with intensity λ, we have $\mathbb{E}[N] = O(\lambda\alpha S^2)$ This yields $O\left((\lambda\alpha S^2)^2\right)$ computational operations each time the planning algorithm is called. We normalize this number by the time it takes the robot to traverse this distance, *i.e.*, by S/v. Then, the asymptotic running time complexity for motion planning is

$$C_P = O\left(\frac{(\lambda\alpha S^2)^2}{S/v}\right) = O\left(\lambda^2\alpha^2 S^3 v\right).$$

b) On inference task: The robot must perform some form of inference each time it visits a target location, processing the data collected at the same location. Hence, the computational complexity of inference tasks is the number of tasks visited. The number of targets visited is analyzed similarly to the amount of reward collected. Following the proof of Theorems7 and 11, we find that the number of targets visited while traversing a distance S is $O(\sqrt{\alpha}\,S)$. We assume that there is a constant number of operations performed at each location for inference. Then, the total number sensing operations is $O(\sqrt{\alpha}\,S)$. We normalize this number with the time it takes to travel distance S to arrive at the number of computational operations per unit time devoted to inference. The time is S/v. Hence, the computational runtime complexity of inference is

$$C_I = O\left(\frac{\sqrt{\alpha}\,S}{S/v}\right) = O\left(\sqrt{\alpha}\,v\right).$$

V. Computational Experiments

This section is devoted to the results of Monte-Carlo simulation studies that verify our theoretical results in Sections III and IV. We consider the discrete and continuous problems separately, and in each case we study the robot performance with different reward distributions, including geometric, exponential, Bernoulli, and Pareto.

A. Optimal Mean Reward on Discrete Lattices

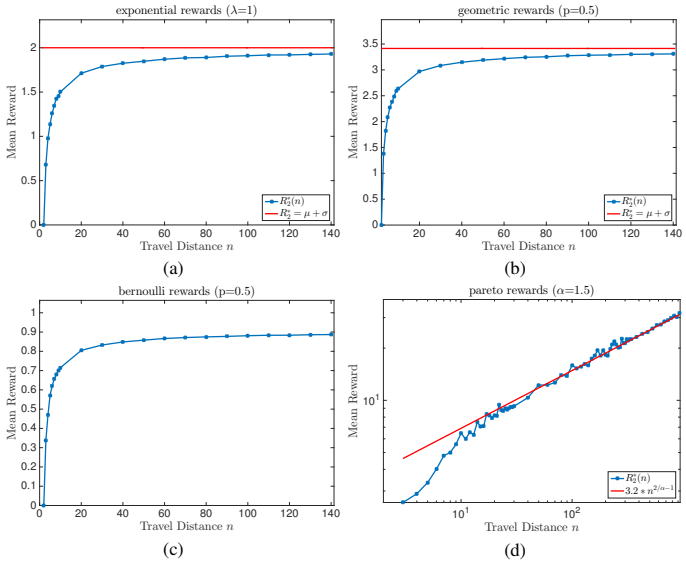

Fig. 3. Optimal mean reward vs. travel distance on a two-dimensional regular lattice. The rewards associated with each vertex on the lattice are independent and identically distributed according to (a) exponential, (b) geometric, (c) Bernoulli and (d) Pareto distributions, respectively. (a)-(c) The mean reward for light-tailed distributions converges to a finite constant. (d) The mean reward for a heavy-tailed distribution goes to infinity.

In this section, we verify Theorems 1-3 and Propositions 1-2 in Monte-Carlo simulations. We consider a robot moving on a two-dimensional regular lattice, where the rewards are independent and identically distributed. We examine the optimal mean reward collected over the course of motion, and the results are shown in Figure 3. In this set of experiments, the rewards follow geometric, exponential, Bernoulli, and Pareto distribution, respectively.

The exponential, geometric, Bernoulli distributions, shown in Figure 3(a)-(c), all belong the light-tailed family. Therefore, as predicted by Theorems 1-3, their mean reward converges quickly towards a finite constant. In addition, the mean reward of both the geometric and exponential distributions converges

to $\mathbf{R}_2^* = \mu+\sigma$ (indicated by the red lines), which is the optimal mean reward predicted by Propositions 1.

On the other hand, Figure 3(d) shows the log-log plot with rewards being Pareto-distributed (with parameter $\alpha = 1.5$). The mean reward increased to infinitey with the travel distance, as dictated by Theorem 3 for all heavy-tailed distributions. Since Figure 3(d) is a log-log plot, the mean reward grows with travel distance n at a rate of $n^{2/\alpha-1}$, as predicted by Proposition 2.

B. Sensing Range on Discrete Lattices

In this section, we verify Theorem 4-5 in Monte-Carlo simulations. We consider a robot moving on a state lattice with limited sensing range m. It runs Algorithm 1. We fix the error level $\delta = 0.1$. We let the robot travel with Algorithm 1 until the mean reward it collects during one iteration of the receding-horizon algorithm is at least δ less than the optimal. In other words, the robot stops if $T_i/m < R_2^* - \delta$ for the i^{th} iteration. In Figure 4, we plot the distance traveled in this way versus the sensing range of the robot. Each data point is an average over 1000 trials. We consider exponential, geometric, Bernoulli and Pareto reward distributions.

Notice that Figure 4(a)-(c) are semi-log plots, so the expected distance of travel scale exponentially with increasing sensing range for the light-tailed distributions. In other words, the sensing range is a logarithm of the distance traveled, as is stated by Theorem 4.

We run a slightly different experiment with Pareto-distributed rewards, where the robot stops if $T_i/m < R_2^*(m^{1.1}) - \delta$ and m is the sensing range. This is due to the fact that, based on Theorem 3, $R_2^* = \infty$ for all heavy-tailed rewards, so the distance-of-travel would have been 0. In Figure 4(d) shows the result, where the distance of travel increases only linearly (not exponentially) and is therefore much worse than the light-tailed distributions. This is consistent with our Theorem 5.

C. Mean Reward in Continuous Spaces

In this section, we verify Theorems 6-8, as well as Proposition 3, in Monte-Carlo simulations. We consider the continuous problem described in Section II-B. The target locations are distributed according to a Poisson process with intensity 1. We consider rewards that are distributed according to the exponential, geometric, Bernoulli and Pareto distributions. In Figure 5, we plot the mean reward versus the travel distance for each of these reward distributions. As predicted by Theorems 6-8, as travel distance increases, the mean reward seems to converge towards a finite value for light-tailed distributions, as shown in Figure 5(a)-(c). In Figure 5(d), however, we show a log-log plot with Pareto rewards. It is clear that the optimal mean reward is not only diverging to infinity, but also growing at a rate of $O(L^{2/\alpha-1})$ as the travel distance L increases, as predicted by Proposition 3.

D. The Impact of Sensing Range On Performance

In this section, we verify Theorem 9-10 in Monte-Carlo simulations. We consider the problem setup presented in Section II-B. The target locations are distributed according to a Poisson process with intensity $\lambda = 1$. The robot has limited sensing range, and it runs Algorithm 2. We fix a sensing range S, and we let the robot travel until the mean reward it collects goes under the value that is δ away from the optimal. We

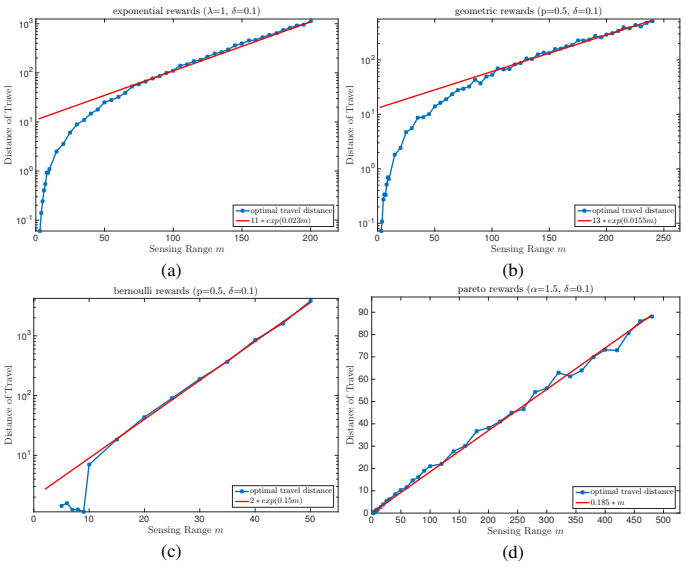

Fig. 4. Average distance-of-travel versus sensing range on a two-dimensional regular lattice. (a)-(c) are log-linear plots for exponential, geometric and bernoulli rewards, respectively. The distance-of-travel grows exponentially fast with sensing range when reward distribution is in the light-tailed family. (d) is a linear plot for pareto rewards (in the heavy-tailed family), and the distance-of-travel only grows linearly with sensing range.

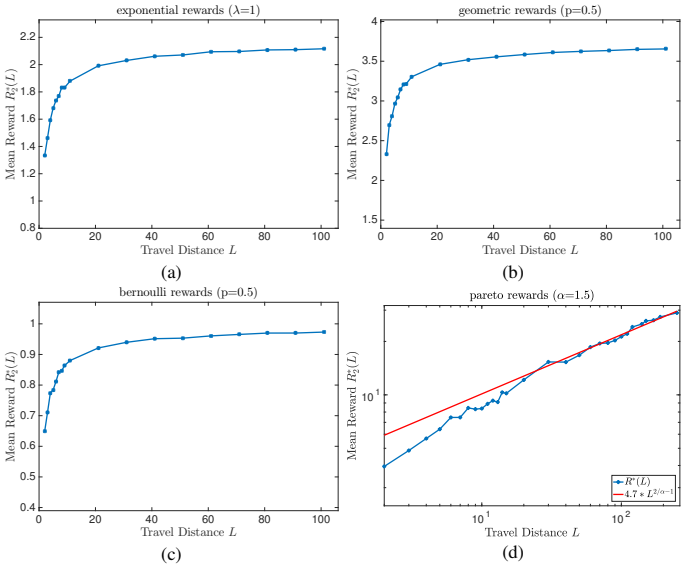

Fig. 5. Optimal mean reward vs. travel distance on a two-dimensional poisson random reward field. The rewards associated with each target in the field are independent and identically distributed according to (a) exponential, (b) geometric, (c) Bernoulli and (d) Pareto distributions, respectively. (a)-(c) The mean reward for light-tailed distributions converges to a finite constant. (d) The mean reward for a heavy-tailed distribution goes to infinity.

record the distance the robot can travel in this manner.

In Figures 6(a)-(c), we plot this distance traveled versus the sensing range in semi-log plots. We started the experiment with for exponentially distributed, geometric distributed and Bernoulli rewards, respectively. Notice that the distance traveled in this manner grows exponentially with increasing sensing range. Hence, in other words, the sensing range required to traverse a certain distance increases only logarithmically with the travel distance. The robot's performance, in terms of the reward collected, is still guaranteed to be a constant factor away from the optimal.

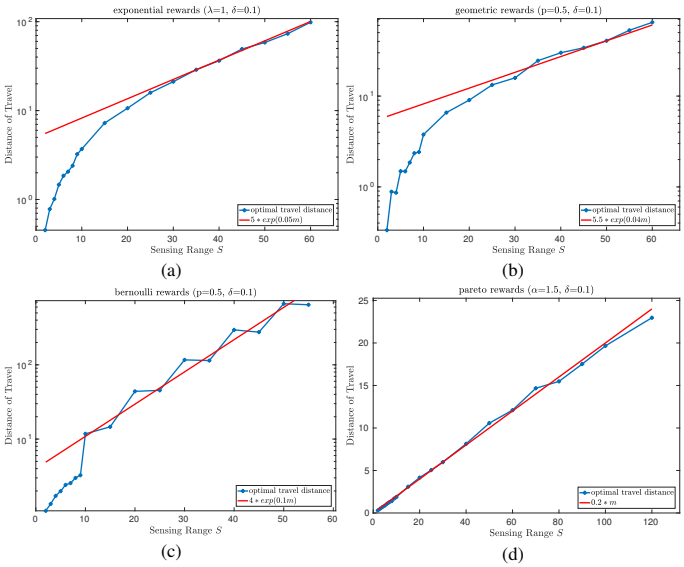

Fig. 6. The average distance-of-travel is plotted against sensing distance S of the robot for the continuous problem. The Poisson process is parameterized with $\lambda = 1$. (a)-(c) are log-linear plots for exponential, geometric and bernoulli rewards, respectively. The distance-of-travel grows exponentially fast with sensing range when reward distribution is in the light-tailed family. (d) is a linear plot for pareto rewards (in the heavy-tailed family), and the distance-of-travel only grows linearly with sensing range.

In Figure 6(d) we show the experiment with Pareto-distributed rewards. Similar to Section V-B, we run a slightly different experiment where the baseline mean reward is $R_2^*(S^{1.1})$ instead of R_2^*, where S is the sensing range. The optimal distance travelled is increasing only linearly with the sensing range S.

E. The Impact of Agility on Performance

In this section, we verify Theorem 11 in Monte-Carlo simulations. We consider the problem setup described in Section II-B. The locations of the targets are distributed according to a Poisson process with intensity $\lambda = 10$, and the rewards are exponentially distributed with unit mean. The travel distance is fixed to $L = 30$. In these simulations, we vary the agility parameter and we observe how the mean reward varies. The results are shown in Figure 7. Each data point is averaged over 300 independent trials.

Notice that the mean reward follows the $\sqrt{A}$ rule, where A is the agility of the robot, as stated by Theorem 11.

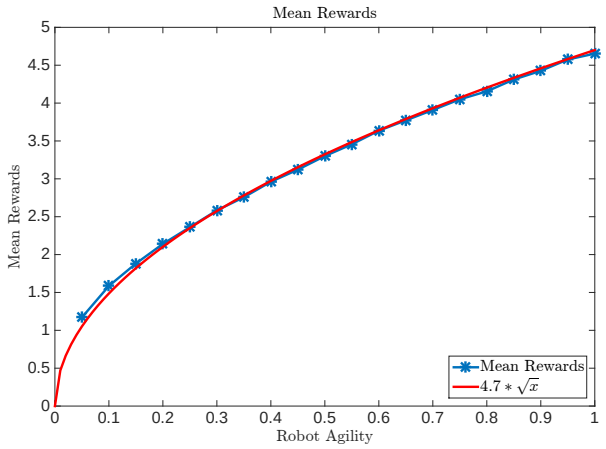

Fig. 7. Mean reward versus robot agility for the simulation where the intensity λ of the Poisson point process is 10 and the reward follows an exponential distribution with mean 1.

F. Requirements on Computation

In this section, we verify our claims in Section IV-E in Monte-Carlo simulations. We consider the setting of the problem presented in Section II-B. The robot travels with limited perception range, and we take a look at how sensing range and agility impacts the computation time devoted to planning and inference tasks. The results are presented in Figure 8 for computation time devoted to motion planning, and in Figure 9 for computation time devoted to inference.

In Figure 8(a), we find that the computation time for motion planning scales as S^4, where S is the sensing range, each time the algorithm is run. Hence, the computation required for motion planning is S^3 per unit distance. In Figure 8(b), we observe that the computation time increases quadratically with increasing values of the agility parameter α. Notice that the computation time devoted to motion planning scales quadratically with increasing agility, as predicted in Section IV-E.

In Figure 9(a), we observe that the computation time devoted to inference (as measured by the number of inference tasks) is constant with increasing sensing range. In Figure 9(b), we see that the computation time devoted to inference grows roughly as $\sqrt{\alpha}$, where α is the agility of the robot, as predicted in Section IV-E.

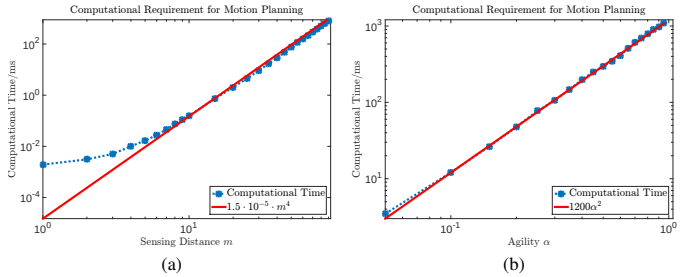

Fig. 8. Mean runtime of dynamic programming for motion planning versus sensing range m and robot agility α. The intensity of the Poisson point process is $\lambda = 1$ and the reward follows a bernoulli distribution with $p = 0.5$.

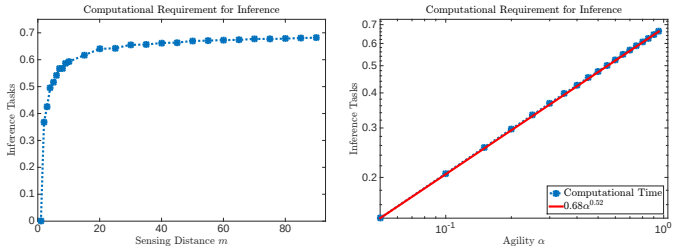

(a) (b)

Fig. 9. Average number of inference tasks (targets visited) versus sensing range m and robot agility α. The number of targets visited grows with the robot's sensing distance but quickly reaches a plateau, and is linear with robot agility α.

VI. Discussion

In this section, we provide a brief discussion on our theoretical results, and illustrate how they can be used to gain valuable insight into design problems involving autonomous vehicles utilized in data gathering applications. First, in Section VI-A, we outline the insights that our major findings presented in Section IV provide. Second, in Section VI-B, we consider a sensor selection problem involving an aerial vehicle gathering data from unattended ground sensors, where we analyze sensing with homogeneous versus heterogeneous sensors.

A. Insight for the Design of Sensing, Agility, and Computation Properties of Data-gathering Vehicles

In this section, we discuss our major results in the context of the data-gathering problem presented in Section II-C from the perspective of sensor precision, sensing range, agility, and computation capabilities. The main objective of this section is to establish the connections between the data-gathering problem presented in Section II-C and the maximum-reward motion problem we presented in II-B and analyzed in Section IV.

a) On sensor precision: Recall from the data-gathering problem that the sensing precision β corresponds to the reward in the maximum-reward motion problem. Hence, the larger the sensor precision β in the former, the higher the corresponding reward in the latter. Then, according to Theorem 6, the precision of the estimate of the unknown variable θ will increase linearly with travel distance L, almost surely. Furthermore, the total precision of the estimate is at least $L \cdot \sqrt{\lambda \mathbb{E}[\beta^2]}$, where λ is the intensity of the target locations and β is a random variable (either geometric or exponential) that denotes the precision of each measurement.

Hence, we find that both increasing measurement precision (β) and increasing travel distance (L) has non-diminishing returns for the data-gathering problem outlined in Section II-C.

b) On sensing range: The sensing range has different implications depending on the distribution of the precision of each measurement. If the precision of the measurements are light-tailed, then the precision of the estimate of θ increases linearly with increasing distance. Furthermore, according to Theorem 9, even with a sensing distance of $\log(L)$, the precision of the estimate of θ is almost as good as the precision of the same estimate when the sensing distance is L, if the vehicle travels a distance of L, with high probability. In other words, it is possible to achieve near-optimal estimation performance with little sensing distance for light-tailed reward.

However, when the precision of the measurements is distributed according to the Pareto distribution with parameter $\alpha \in (0, 2)$, then the precision of the estimate increases super-linearly with increasing sensing distance. Furthermore, according to Theorem 10, it is impossible to obtain near-optimal estimation performance with small sensing distance for Pareto reward with parameter $\alpha \in (0, 2)$, which is a heavy-tailed distribution. We conjecture that this result applies to all heavy tailed distributions of the precision of the measurements.

c) On agility: According to Theorem 11, the precision of the estimate increases with increasing agility $\alpha = w/v$, where w is the maximum lateral speed and v is the longitudinal speed of the vehicle (see Section II-B). However, the increase comes with diminishing returns, proportional to $\sqrt{\alpha}$.

d) On computation workload: The computational workload is determined with the sensing range and the agility of the robot. The quantification of the computational workload for the data-gathering problem of Section II-C follows that of the maximum-reward problem, the analysis for which was presented in Section IV-E.

The computational workload can be partitioned into two activities, namely motion planning and inference. The motion planning task consists of determining the set of target locations to be visited each time a new target gets in the sensing distance of the vehicle. The computational workload for this task increases substantially with increasing sensing distance and robot agility. Specifically, the the computational workload for planning increases as $O(\alpha^2 S^3)$ per unit time, where S is the sensing distance and α is the robot agility.

The inference task consists of incorporating the new measurements to improve the estimate. This task may be computationally challenging as it may involve image analysis, sensor fusion, *et cetera*. The computational workload for this task increases proportionally with $\sqrt{\alpha}$, where α is the agility of the robot. It is independent of the sensing range.

B. Case Study: Unattended Ground Sensor Selection

In this section, we present a short case study that involves a UAV is tasked with estimating an unknown variable θ with data acquired from Unattended Ground Sensors (UGS) that are randomly distributed over a region of interest.

The UGS technology is an emerging technology that may have substantial impact in environmental monitoring, surveillance, and reconnaissance. The UGS often house primitive sensors that record various measurements, *e.g.*, seismic, acoustic, magnetic, temperature, and humidity measurements, continuously for extended time periods, *e.g.*, for several months. They are often deployed sparsely, which prevents formation of ad-hoc networks. However, the data they record can be collected by UAVs that fly over the sensors.

In this section, we demonstrate how our analysis can be utilized to arrive at fundamental results for a certain kind of UGS selection problem. Specifically, we consider a problem

where each UGS provides a measurement of the hidden variable θ corrupted with Gaussian noise. The precision of the measurement may depend on the quality of the UGS. We assume that the UAV recognizes each UGS from a certain distance, and learns the precision of the measurement that is obtained by that UGS. The UAV must plan its path carefully to best estimate the unknown variable θ. Clearly, this problem is the same as the problem which we presented in Section II-C.

Suppose we have the option to choose the sensors before they are distributed in the field. Due to limited budgets, the average quality of sensors is fixed and the total number of sensors is given, *i.e.*,

$$\mathbb{E}[\beta_i] = \mu_\beta,$$

where μ_β is some positive constant and λ is known. With above constraints, we would like to address the following question: *Which one of the following two strategies yields a higher level of confidence for the estimation?*

1) *Assign the same level of precision to all sensors,* i.e., $\beta_i = \mu_\beta$ *for all* i
2) *Randomize the level of precision* β_i *over some probability distributions* F_β *with mean* μ_β

In the former option, all sensors are of the same quality; in the latter one, some sensors provide more precise (less noisy) measurements, while some others provide less provide (more noisy) measurements, when compared to the former option.

Notice that this sensor selection problem is an instance of the maximum-reward motion problem presented in Section II-B. In this case, the reward is the precision β_i of the measurement that the ith UGS provides, hence the quality of that UGS.

Let us analyze the performance of each of the strategies by computing the total reward collected in each case. As we established in the previous section, the total reward is proportional to the precision of our estimate of the hidden variable θ by visiting the unattended ground sensors.

The first strategy assigns equal precision to all UGS sensors. For optimal performance, the robot should visit as many sensors as possible in order to maximize the total precision gain. The resulting performance is analyzed below.

Theorem 12 (Adapted From [62])**.** *Suppose the reward locations are generated by a Poisson point process with intensity* λ *on* $\mathbb{R}^2$ *and all reward is 1. The robot dynamics satisfies the following ordinary differential equation:*

$$\dot{x}_1(t) = v, \quad \dot{x}_2(t) = u(t),$$

where $|u(t)| \leq v$ (i.e.*, the robot agility is 1). Then*

$$\lim_{L\to\infty} \frac{T(L)}{L} = \sqrt{2\lambda} \quad \textit{almost surely}.$$

Theorem 12 provides the expected optimal mean reward collected when the rewards are equal to 1 surely. Therefore, it follows that the average number of sensors visited is $\sqrt{2\lambda}$, and hence the overall precision gain would be $L \cdot \sqrt{2\lambda} \cdot \mathbb{E}[\beta]$.

The second strategy, on the other hand, utilizes sensors with random precisions (for instance, when the precisions follow an exponential distribution with mean equal to the homogeneous strategy). The exact R_2^* for the Poisson random reward field is out of reach at the moment, but the computational experiments in Figure 4(a) show that when the mean $\mathbb{E}[\beta] = 1$, the expected precision gain is at least 2.1, much higher than the expected precision gain $\sqrt{2}$ when using homogeneous sensors. By the comparison of light-tailed (bounded variance) and heavy-tailed (infinite variance) distributions, we conjecture that the higher variance of the distributions, the better performance we can expect from the deployment of random-precision sensors.

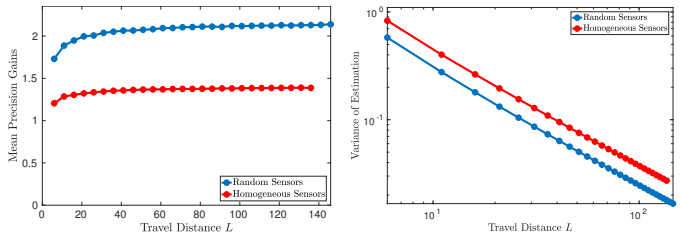

(a) (b)

Fig. 10. Comparison of randomized sensors against homogeneous sensors by the mean precision gains and the variance of estimation. The randomized strategy assumes that sensor precisions follow an exponential distribution with mean 1, while the homogeneous sensors have a constant precision of 1. The randomized strategy provides higher mean precision gain with lower variance.

To illustrate this comparison, we present results of a Monte-Carlo study. In Figure 10(a), we show the precision gains, and in Figure 10(b) we show how the variance of estimate (*i.e.*, inverse of the precision gains) of the hidden variable decays. We observe that, when the sensor quality is randomized, the quality of the estimate is better and the variance of estimation decreases faster.

VII. Conclusion

In this paper, we propose the maximum-reward motion problem for studying fundamental limits of data-gathering robots, given their sensing, actuation and computation constraints. We model the robot as a particle moving in a stochastic reward field and analyze its performance by using results from last-passage percolation problem in statistical mechanics. We verify our theoretical results in thorough simulation experiments. We also apply our results in the design of data-gathering vehicles as well as sensor selection for unattended ground sensors, providing insights for these problems.

References

[1] M. Bellmore and G. L. Nemhauser, "The traveling salesman problem: a survey," *Operations Research*, vol. 16, no. 3, pp. 538–558, 1968.

[2] E. L. Lawler, "The traveling salesman problem: a guided tour of combinatorial optimization," *WILEY-INTERSCIENCE SERIES IN DISCRETE MATHEMATICS*, 1985.

[3] G. Laporte, "The traveling salesman problem: An overview of exact and approximate algorithms," *European Journal of Operational Research*, vol. 59, no. 2, pp. 231–247, 1992.

[4] D. L. Applegate, R. E. Bixby, V. Chvatal, and W. J. Cook, *The Traveling Salesman Problem: A Computational Study*. Princeton university press, 2011.

[5] S. Lin and B. W. Kernighan, "An effective heuristic algorithm for the traveling-salesman problem," *Operations research*, vol. 21, no. 2, pp. 498–516, 1973.

[6] D. J. Rosenkrantz, R. E. Stearns, and P. Lewis, "Approximate algorithms for the traveling salesperson problem," in *Switching and Automata Theory, 1974., IEEE Conference Record of 15th Annual Symposium on*. IEEE, 1974, pp. 33–42.

[7] M. Padberg and G. Rinaldi, "A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems," *SIAM review*, vol. 33, no. 1, pp. 60–100, 1991.

[8] M. Dorigo and L. Gambardella, "Ant-q: A reinforcement learning approach to the traveling salesman problem," in *Proceedings of ML-95, Twelfth Intern. Conf. on Machine Learning*, 2014, pp. 252–260.

[9] J. Beardwood, J. H. Halton, and J. M. Hammersley, "The shortest path through many points," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 55, no. 04. Cambridge Univ Press, 1959, pp. 299–327.

[10] K. Savla, E. Frazzoli, and F. Bullo, "Traveling salesperson problems for the dubins vehicle," *Automatic Control, IEEE Transactions on*, vol. 53, no. 6, pp. 1378–1391, 2008.

[11] J. Le Ny, E. Feron, and E. Frazzoli, "On the dubins traveling salesman problem." *IEEE Trans. Automat. Contr.*, vol. 57, no. 1, pp. 265–270, 2012.

[12] G. B. Dantzig and J. H. Ramser, "The truck dispatching problem," *Management science*, vol. 6, no. 1, pp. 80–91, 1959.

[13] N. Christofides, "The vehicle routing problem," *Revue française d'automatique, d'informatique et de recherche opérationnelle. Recherche opérationnelle*, vol. 10, no. 1, pp. 55–70, 1976.

[14] P. Toth and D. Vigo, *The vehicle routing problem*. Society for Industrial and Applied Mathematics, 2001.

[15] B. L. Golden, S. Raghavan, and E. A. Wasil, *The Vehicle Routing Problem: Latest Advances and New Challenges: latest advances and new challenges*. Springer Science & Business Media, 2008, vol. 43.

[16] G. Laporte, "The vehicle routing problem: An overview of exact and approximate algorithms," *European Journal of Operational Research*, vol. 59, no. 3, pp. 345–358, 1992.

[17] M. Desrochers, J. Desrosiers, and M. Solomon, "A new optimization algorithm for the vehicle routing problem with time windows," *Operations research*, vol. 40, no. 2, pp. 342–354, 1992.

[18] I. H. Osman, "Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem," *Annals of operations research*, vol. 41, no. 4, pp. 421–451, 1993.

[19] M. Gendreau, A. Hertz, and G. Laporte, "A tabu search heuristic for the vehicle routing problem," *Management science*, vol. 40, no. 10, pp. 1276–1290, 1994.

[20] R. Baldacci, A. Mingozzi, and R. Roberti, "Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints," *European Journal of Operational Research*, vol. 218, no. 1, pp. 1–6, 2012.

[21] N. H. Wilson and N. J. Colvin, *Computer control of the Rochester dial-a-ride system*. Massachusetts Institute of Technology, Center for Transportation Studies, 1977.

[22] M. Gendreau, G. Laporte, and R. Séguin, "Stochastic vehicle routing," *European Journal of Operational Research*, vol. 88, no. 1, pp. 3–12, 1996.

[23] D. J. Bertsimas and G. Van Ryzin, "A stochastic and dynamic vehicle routing problem in the euclidean plane," *Operations Research*, vol. 39, no. 4, pp. 601–615, 1991.

[24] S. Kataoka and S. Morito, "An algorithm for single constraint maximum collection problem." *J. OPER. RES. SOC. JAPAN.*, vol. 31, no. 4, pp. 515–530, 1988.

[25] S. E. Butt and T. M. Cavalier, "A heuristic for the multiple tour maximum collection problem," *Computers & Operations Research*, vol. 21, no. 1, pp. 101–111, 1994.

[26] G. Laporte and S. Martello, "The selective travelling salesman problem," *Discrete applied mathematics*, vol. 26, no. 2, pp. 193–207, 1990.

[27] D. Feillet, P. Dejax, and M. Gendreau, "Traveling salesman problems with profits," *Transportation science*, vol. 39, no. 2, pp. 188–205, 2005.

[28] E. M. Arkin, J. S. Mitchell, and G. Narasimhan, "Resource-constrained geometric network optimization," in *Proceedings of the fourteenth annual symposium on Computational geometry*. ACM, 1998, pp. 307–316.

[29] B. L. Golden, L. Levy, and R. Vohra, "The orienteering problem," *Naval Research Logistics (NRL)*, vol. 34, no. 3, pp. 307–318, 1987.

[30] P. Vansteenwegen, W. Souffriau, and D. Van Oudheusden, "The orienteering problem: A survey," *European Journal of Operational Research*, vol. 209, no. 1, pp. 1–10, 2011.

[31] T. Tsiligirides, "Heuristic methods applied to orienteering," *Journal of the Operational Research Society*, pp. 797–809, 1984.

[32] I.-M. Chao, B. L. Golden, and E. A. Wasil, "A fast and effective heuristic for the orienteering problem," *European Journal of Operational Research*, vol. 88, no. 3, pp. 475–489, 1996.

[33] B. L. Golden, Q. Wang, and L. Liu, "A multifaceted heuristic for the orienteering problem," *Naval Research Logistics (NRL)*, vol. 35, no. 3, pp. 359–366, 1988.

[34] M. Fischetti, J. J. S. Gonzalez, and P. Toth, "Solving the orienteering problem through branch-and-cut," *INFORMS Journal on Computing*, vol. 10, no. 2, pp. 133–148, 1998.

[35] A. C. Leifer and M. B. Rosenwein, "Strong linear programming relaxations for the orienteering problem," *European Journal of Operational Research*, vol. 73, no. 3, pp. 517–523, 1994.

[36] Z. W. Geem, C.-L. Tseng, and Y. Park, "Harmony search for generalized orienteering problem: best touring in china," in *Advances in natural computation*. Springer, 2005, pp. 741–750.

[37] I.-M. Chao, B. L. Golden, and E. A. Wasil, "The team orienteering problem," *European journal of operational research*, vol. 88, no. 3, pp. 464–474, 1996.

[38] H. Tang and E. Miller-Hooks, "A tabu search heuristic for the team orienteering problem," *Computers & Operations Research*, vol. 32, no. 6, pp. 1379–1407, 2005.

[39] C. Archetti, A. Hertz, and M. G. Speranza, "Metaheuristics for the team orienteering problem," *Journal of Heuristics*, vol. 13, no. 1, pp. 49–76, 2007.

[40] P. Vansteenwegen, W. Souffriau, G. V. Berghe, and D. Van Oudheusden, "Iterated local search for the team orienteering problem with time windows," *Computers & Operations Research*, vol. 36, no. 12, pp. 3281–3290, 2009.

[41] ——, "A guided local search metaheuristic for the team orienteering problem," *European Journal of Operational Research*, vol. 196, no. 1, pp. 118–127, 2009.

[42] N. Labadie, R. Mansini, J. Melechovskỳ, and R. W. Calvo, "The team orienteering problem with time windows: An lp-based granular variable neighborhood search," *European Journal of Operational Research*, vol. 220, no. 1, pp. 15–27, 2012.

[43] R. N. Smith, M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme, "Persistent ocean monitoring with underwater gliders: Adapting sampling resolution," *Journal of Field Robotics*, vol. 28, no. 5, pp. 714–741, 2011.

[44] S. Garg and N. Ayanian, "Persistent monitoring of stochastic spatio-temporal phenomena with a small team of robots," *Proceedings of Robotics: Science and Systems, Berkeley, USA*, 2014.

[45] C. G. Cassandras, X. C. Ding, and X. Lin, "An optimal control approach for the persistent monitoring problem," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*. IEEE, 2011, pp. 2907–2912.

[46] X. Lin and C. Cassandras, "An optimal control approach to the multi-agent persistent monitoring problem in two-dimensional spaces," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*. IEEE, 2013, pp. 6886–6891.

[47] C. G. Cassandras, X. Lin, and X. Ding, "An optimal control approach to the multi-agent persistent monitoring problem," *Automatic Control, IEEE Transactions on*, vol. 58, no. 4, pp. 947–961, 2013.

[48] S. L. Smith, M. Schwager, and D. Rus, "Persistent monitoring of changing environments using a robot with limited range sensing," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 5448–5455.

[49] ——, "Persistent robotic tasks: Monitoring and sweeping in changing environments," *Robotics, IEEE Transactions on*, vol. 28, no. 2, pp. 410–426, 2012.

[50] S. Alamdari, E. Fata, and S. L. Smith, "Persistent monitoring in discrete environments: Minimizing the maximum weighted latency between observations," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 138–154, 2014.

[51] ——, "Min-max latency walks: Approximation algorithms for monitoring vertex-weighted graphs," in *Algorithmic Foundations of Robotics X*. Springer, 2013, pp. 139–155.

[52] J. Yu, S. Karaman, and D. Rus, "Persistent monitoring of events with stochastic arrivals at multiple stations," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 5758–5765.

[53] S. L. Smith, S. D. Bopardikar, and F. Bullo, "A dynamic boundary guarding problem with translating targets," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*. IEEE, 2009, pp. 8543–8548.

[54] S. D. Bopardikar, S. L. Smith, F. Bullo, and J. P. Hespanha, "Dynamic vehicle routing for translating demands: Stability analysis and receding-horizon policies," *Automatic Control, IEEE Transactions on*, vol. 55, no. 11, pp. 2554–2569, 2010.

[55] N. Matni and V. Chandrasekaran, "Regularization for design," California Institute of Technology, Tech. Rep., 2015.

[56] A. Censi, "A mathematical theory of co-design," Laboratory for Information and Decision Systems/MIT, Tech. Rep., January 2016.

[57] ——, "Uncertainty in monotone co-design problems," *IEEE Robotics and Automation Letters*, 2017.

[58] S. Karaman and E. Frazzoli, "High-speed flight in an ergodic forest," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 2899–2906.

[59] ——, "High-speed motion with limited sensing range in a poisson forest," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*. IEEE, 2012, pp. 3735–3740.

[60] S. Choudhury, S. Scherer, and J. A. Bagnell, "Theoretical limits of speed and resolution for kinodynamic planning in a poisson forest."

[61] L. T. Rolla and A. Q. Teixeira, "Last passage percolation in macroscopically inhomogeneous media," *Electronic Communications in Probability*, vol. 13, pp. 131–139, 2008.

[62] T. Seppäläinen *et al.*, "Increasing sequences of independent points on the planar lattice," *The Annals of Applied Probability*, vol. 7, no. 4, pp. 886–898, 1997.

[63] T. Seppäläinen, "Lecture notes on the corner growth model," *Unpublished notes*, 2009.

[64] X. Zeng, Z. Hou, C. Guo, and Y. Guo, "Directed last-passage percolation and random matrices," *Advances in Mathematics*, vol. 42, no. 3, p. 3, 2013.

[65] G. Grimmet and P. Heimer, "Directed Percolation and Random Walk," in *In and out of equilibrium: probability with a physics flavor*. Birkhauser, 2002.

[66] B. Hambly and J. B. Martin, "Heavy tails in last-passage percolation," *Probability Theory and Related Fields*, 2007.

[67] F. Baccelli, A. Borovkov, and J. Mairesse, "Asymptotic results on infinite tandem queueing networks," *Probability theory and related fields*, vol. 118, no. 3, pp. 365–405, 2000.

[68] P. W. Glynn and W. Whitt, "Departures from many queues in series," *The Annals of Applied Probability*, pp. 546–572, 1991.

[69] J. B. Martin, "Last-passage percolation with general weight distribution," *Markov Process. Related Fields*, vol. 12, no. 2, pp. 273–299, 2006.

[70] A. Somanath, S. Karaman, and K. Youcef-Toumi, "Controlling stochastic growth processes on lattices: Wildfire management with robotic fire extinguishers," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 1432–1437.

[71] F. Ma and S. Karaman, "Maximum-reward motion in a stochastic environment: The nonequilibrium statistical mechanics perspective," in *Algorithmic Foundations of Robotics XI*. Springer, 2015, pp. 389–406.

[72] V. G. I. T. T. Ivancevic, *Applied differential geometry: a modern introduction*. World Scientific, 2007.

[73] D. J. Daley and D. Vere-Jones, *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media, 2007, vol. 2.

[74] S. N. Chiu, D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic geometry and its applications*. John Wiley & Sons, 2013.

[75] H. J. Kushner, *Introduction to stochastic control*. Holt, Rinehart and Winston New York, 1971.

[76] C. U. et al., "Autonomous driving in urban environments: Boss and the Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.

[77] S. Koenig, M. Likhachev, and D. Furcy, "Lifelong planning A$*$," *Artificial Intelligence*, 2004.

[78] M. Pivtoraiko and A. Kelly, "Kinodynamic motion planning with state lattice motion primitives," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*. IEEE, 2011, pp. 2172–2179.

[79] M. Cirillo, T. Uras, and S. Koenig, "A lattice-based approach to multi-robot motion planning for non-holonomic vehicles," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. IEEE, 2014, pp. 232–239.

[80] A. Schrijver, *Combinatorial optimization*. Springer Verlag, 2003.

[81] T. Rolski, H. Schmidli, V. Schmidt, and J. Teugels, *Stochastic processes for insurance and finance*. John Wiley & Sons, 2009, vol. 505.

[82] O. Angel and A. Tomberg, "Last passage percolation," *Mprime Summer School in Probability*, 2012. [Online]. Available: http://www.math.ubc.ca/~angel/ssprob12/courses.php

[83] E. Cator and L. P. Pimentel, "Hydrodynamical methods in last passage percolation models," *arXiv preprint arXiv:1106.2687*, 2011.

[84] M. J. Steele, *Probability Theory and Combinatorial Optimization*. SIAM, 1996.

[85] K. Johansson, "Random growth and random matrices," in *European Congress of Mathematics*. Springer, 2001, pp. 445–456.

[86] B. Hambly and J. B. Martin, "Heavy tails in last-passage percolation," *probability theory and related fields*, vol. 137, no. 1-2, pp. 227–275, 2020.

[87] P. Billingsley, *Convergence of probability measures*. John Wiley & Sons, 2009, vol. 493.

[88] H. L. Royden and P. Fitzpatrick, *Real analysis*. Macmillan New York, 1988, vol. 198, no. 8.

[89] J. B. Martin, "Linear growth for greedy lattice animals," *Stochastic Processes and their Applications*, vol. 98, no. 1, pp. 43–66, 2002.

[90] M. Talagrand, "Concentration of measure and isoperimetric inequalities in product spaces," *Publications Mathématiques de l'Institut des Hautes Etudes Scientifiques*, vol. 81, no. 1, pp. 73–205, 1995.

[91] E. Hille and R. S. Phillips, *Functional analysis and semi-groups*. American Mathematical Soc., 1996, vol. 31.

[92] R. L. Streit, "The poisson point process," in *Poisson Point Processes*. Springer, 2010, pp. 11–55.

APPENDIX A
PROOF FOR THEOREM 1

Let's first introduce the definition of *subadditivity* and the *Fekete's Subadditive Lemma*.

Definition 3 (Subadditivity [84])**.** *A sequence* $\{a_n\}, n \geq 1$*, is called* subadditive *if it satisfies the inequality*

$$a_{n+m} \leq a_n + a_m.$$

A sequence is superadditive *if*

$$a_{n+m} \geq a_n + a_m.$$

Lemma 1 (Fekete's Subadditive Lemma [84])**.** *For every subadditive sequence* $\{a_n\}_{n=1}^{\infty}$*, the limit* $\lim_{n\to\infty} \frac{a_n}{n}$ *exists and is equal to* $\inf \frac{a_n}{n}$.

Now we return to the proof for Theorem 1.

Proof: From Proposition 2.1 in [69] we learn that $\mathbb{E}[T_d^*(n)]$ is superadditive, i.e., the sequence $-\mathbb{E}[T_d^*(n)]$ is subadditive. Then the result follows directly from Lemma 1. ∎

APPENDIX B
PROOF FOR THEOREM 2

The problem of computing the maximum reward that the optimal path on a lattice can possibly collect is called the *last passage percolation problem* [69] in statistical mechanics, and it has connections with non-equilibrium statistical mechanics problems involving corner growth [85]. This problem has attracted tremendous attention in the past. In particular, the function $T_d^*(v_{\text{dest}})$ has been analyzed extensively. Let us recall some of the main results from this literature and subsequently state and prove our main results for this section.

Definition 4 (Shape function)**.** *The function* $g(\cdot)$*, defined as*

$$g(\mathbf{v}) := \sup_{k\in\mathbb{N}} \frac{\mathbb{E}[T_d^*(\lfloor k\,\mathbf{v}\rfloor)]}{k},$$

is called the shape function.

Firstly, the maximum reward towards any destination v has been shown to converge to a limit, when reward is independent and identically distributed.

Proposition 4 (See Proposition 2.1 in [69])**.** *Assume the reward* $r(v)$ *at each vertex* $\mathbf{v}$ *is an i.i.d. random variable and* $\mathbb{E}[r(\mathbf{v})] < \infty$*. Then,* $\frac{T_d^*(\lfloor k\,\mathbf{v}\rfloor)}{k}$ *converges to the shape function* $g(\mathbf{v})$ *almost surely as* k *diverges to infinity,* i.e.*,*

$$\mathbb{P}\Big(\lim_{k\to\infty} \frac{T_d^*(\lfloor k\,\mathbf{v}\rfloor)}{k} = g(\mathbf{v})\Big) = 1,$$

and for all $\alpha > 0$*,*

$$\alpha g(\mathbf{v}) = g(\alpha\mathbf{v}).$$

That is, for any $\mathbf{v}$, the maximum reward $T_d^*(\mathbf{v})$ converges to the shape function $g(\mathbf{v})$, almost surely. However, this limit may be infinite, depending on the distribution of $r(\mathbf{v})$. Martin [69] provides a almost optimal necessary condition for the finiteness of the shape function g as follows.

Theorem 13 (See Theorem 4.1 in [69])**.** *If the reward distribution* F *satisfies Equation 5, i.e.,*

$$\int_0^{\infty} (1 - F(s))^{1/d} ds < \infty,$$

then $g(\mathbf{v}) < \infty$ *for all* $\mathbf{v} \in \mathbb{R}_+^d$.

Now we proceed to the proof for Theorem 2.

Proof: By the superadditivity of $\mathbb{E}[T_d^*(n)]$ we have

$$\mathbb{E}[T_d^*(n\mathbf{1})] \leq \mathbb{E}[T_d^*(nd)] \leq \mathbb{E}[T_d^*(2nd\mathbf{1})/2],$$

where $\mathbf{1} = (1, 1, \ldots, 1)$. If we divide the above inequalities all by nd, then for large n both the left-hand-side and right-hand-side converge to the same constant due to the convergence result in Proposition 4. It immediately follows that

$$\lim_{n\to\infty} \mathbb{E}\left[\frac{T_d^*(nd)}{nd}\right] = \lim_{n\to\infty} \mathbb{E}\left[\frac{T_d^*(n\mathbf{1})}{nd}\right] = \frac{g(\mathbf{1})}{d}.$$

Therefore, if Equation 5 holds, we have

$$\mathbf{R}_d^* = \lim_{n\to\infty} \mathbb{E}[R_d^*(n)] = \lim_{n\to\infty} \mathbb{E}\left[\frac{T_d^*(n)}{n}\right] = \frac{g(\mathbf{1})}{d}, \quad (7)$$

which is a finite constant by Theorem 13. ∎

APPENDIX C
PROOF FOR PROPOSITION 1

The results in [69] show that the shape function $g(v)$ can be computed exactly for at least two cases, namely, when the distribution F of reward $r(v)$ at each vertex is either an exponential distribution or a geometric distribution. More specifically, if the reward distribution F is an exponential distribution with parameter $\lambda = 1$, then the shape function g defined in Proposition 4 is

$$g\big((v_1, v_2)\big) = (\sqrt{v_1} + \sqrt{v_2})^2, \quad \text{for all } (v_1, v_2) \in \mathbb{N}^2. \quad (8)$$

If the reward distribution F is a geometric distribution with parameter p, *i.e.*, $\mathbb{P}(X = k) = p(1-p)^{k-1}$ for $k = 1, 2, \ldots,$ then the shape function is

$$g\big((v_1, v_2)\big) = \frac{v_1 + 2\sqrt{v_1 v_2(1-p)} + v_2}{p}, \text{ for all } (v_1, v_2) \in \mathbb{N}^2.$$

Following Equation 7 in Appendix B, we readily derive that for both the geometric and exponential distributions

$$\mathbf{R}_2^* = g(\mathbf{1})/2 = \mu + \sigma.$$

Computing the shape function for other reward distributions, however, remains a long-standing, well-known open problem [69].

APPENDIX D
PROOF FOR PROPOSITION 2

Let's first introduce a lemma that is useful for our proof.

Lemma 2 (See Theorem 2.1 in [86])**.** *Suppose the CDF $F(x)$ is regularly varying with index $\alpha \in (0,2)$. Let $a_N = F^{-1}(1 - 1/N)$, for all $N \in \mathbb{N}$. Then, $a_{n^2}^{-1} T(n)$ converges in distribution to a random variable T that is almost surely finite.*

Now, we return to the proof of Proposition 2. For the Pareto distribution with parameters $x_m > 0$ and $\alpha \in (0,1)$, it is easy to derive that

$$a_{n^2} = F^{-1}(1 - 1/n^2) = x_m \cdot n^{2/\alpha}$$

By Lemma 2, we have that $a_{n^2}^{-1} T_2^*(n)$ converges to a (non-trivial) random variable T in distribution with $T < \infty$ almost surely. That is,

$$\frac{T_2^*(n)}{n^{2/\alpha}} \to x_m \cdot T \quad \text{in distribution} \tag{9}$$

Note that $x_m \cdot T$ is almost surely bounded, which implies a finite expectation. It follows by the definition of R_2^* that

$$\frac{R_2^*(n)}{n^{2/\alpha - 1}} \to x_m \cdot T \quad \text{in distribution.}$$

From Theorem 1 we know that $R_2^*(n) \to R_2^*$ surely. Therefore, it can be seen that as depending on the growth rate of $R_2^*(n)$ with respect to n, $\frac{R_2^*(n)}{n^{2/\alpha-1}}$ might converge to 0, a finite positive constant c, or ∞.

By Skorokhod's Representation Theorem [87], there exists a sequence of random variables X_n, defined on the same probability space, such that X_n has the same distribution as $\frac{R_2^*(n)}{n^{2/\alpha-1}}$ and that $X_n \to x_m \cdot T$ almost surely. Therefore we can apply Fatou's Lemma [88] here and obtain

$$\liminf_{n\to\infty} \mathbb{E} X_n \geq \mathbb{E}\left[x_m \cdot T\right] > 0$$

which eliminates the possibility that X_n (and therefore $\frac{R_2^*(n)}{n^{2/\alpha-1}}$) converges to 0. In other words, $\frac{R_2^*(n)}{n^{2/\alpha-1}}$ must converge to either a finite constant $c > 0$ or ∞, and hence $R_2^*(n)$ is at least of growth rate $n^{2/\alpha-1}$. It immediately follows that

$$R_2^* = \lim_{n\to\infty} \mathbb{E}[R_2^*(n)] = \infty.$$

APPENDIX E
PROOF FOR THEOREM 4

To prove Theorem 4, we adopt a two-step strategy: we first "truncate" the original problem with light-tailed reward to a simplified problem with bounded reward, and then apply concentration of measure to show that the difference between the two problems is actually small.

A. "Truncated" Problem with Bounded Reward

Now let's first consider a truncated version of the original problem. Instead of having light-tailed distributed reward, we define a new problem where the reward associated with vertex $\mathbf{v}$ is now defined as $r^{(L)} = \max(r(\mathbf{v}), L)$, where $L > 0$. We state the result for the truncated problem as follows.

Lemma 3. *Suppose the rewards $r(v)$ are independent, identically distributed and almost surely bounded. Then, for any $\delta > 0$, there exists a constant c such that $R_{\text{iterative}}(n, c\log n)$ converges to $\mathbf{R}_d^*$ in probability,* i.e.*,*

$$\lim_{n\to\infty} \mathbb{P}\left(\left| R_{\text{iterative}}(n, c\log n) - \mathbf{R}_d^* \right| \geq \delta \right) = 0.$$

Before proving Lemma 3, we state an intermediate result that enables our proof. This intermediate result is a concentration inequality, which plays a key role in deriving many results in nonequilibrium statistical mechanics [89].

Lemma 4 (See [89])**.** *Let $\{Y_i, i \in \mathcal{I}\}$ be a finite collection of independent random variables that are bounded almost surely,* i.e.*, $\mathbb{P}(|Y_i| \leq L) = 1$ for all $i \in \mathcal{I}$. Let $\mathcal{C}$ be a collection of subsets of $\mathcal{I}$ with maximum cardinality R,* i.e.*, $\max_{C\in\mathcal{C}} |C| \leq R$ and let $Z = \max_{C\in\mathcal{C}} \sum_{i\in C} Y_i$. Then for any $u > 0$,*

$$\mathbb{P}(|Z - \mathbb{E}Z| \geq u) \leq \exp\left(-\frac{u^2}{64RL^2} + 64 \right).$$

Now, we present the proof for Lemma 3.

Proof for Lemma 3: First, note that:

$$\begin{aligned}
&\lim_{n\to\infty} \mathbb{P}\left(\left|R_{\text{iterative}}(n,m) - R_d^*\right| \geq \delta\right) \\
&= \lim_{n\to\infty} \mathbb{P}\left(\left|R_{\text{iterative}}(n,m) - \frac{\mathbb{E}[T(n)]}{n} + \frac{\mathbb{E}[T(n)]}{n} - R_d^*\right| \geq \delta\right) \\
&= \lim_{n\to\infty} \mathbb{P}\left(\left|R_{\text{iterative}}(n,m) - \frac{\mathbb{E}[T(n)]}{n}\right| \geq \frac{\delta}{2}\right) \\
&\qquad + \lim_{n\to\infty} \mathbb{P}\left(\left|\frac{\mathbb{E}[T(n)]}{n} - R_d^*\right| \geq \frac{\delta}{2}\right)
\end{aligned}$$

The second term is 0 by Theorem 1, so we focus on the first

term. Let's define $\delta = \frac{u}{n}$. By the definition of $R_{\text{iterative}}(n,m)$:

$$\lim_{n\to\infty} \mathbb{P}\left(\left|R_{\text{iterative}}(n,m) - R_d^*\right| \geq \delta\right)$$
$$= \lim_{n\to\infty} \mathbb{P}\left(\left|R_{\text{iterative}}(n,m) - \frac{\mathbb{E}T(n)}{n}\right| \geq \frac{\delta}{2}\right) \tag{10}$$
$$= \lim_{n\to\infty} \mathbb{P}\left(\left|\frac{\sum_{i=1}^{\frac{n}{m}} T_i(m)}{n} - \frac{\mathbb{E}T(n)}{n}\right| \geq \frac{\delta}{2}\right)$$
$$= \lim_{n\to\infty} \mathbb{P}\left(\left|\frac{\sum_{i=1}^{\frac{n}{m}} T_i(m)}{n} - \frac{\mathbb{E}T(m)}{m} + \frac{\mathbb{E}T(m)}{m} - \frac{\mathbb{E}T(n)}{n}\right| \geq \frac{\delta}{2}\right) \tag{11}$$
$$\leq \lim_{n\to\infty} \mathbb{P}\Bigg(\left\{\left|\frac{\sum_{i=1}^{\frac{n}{m}} T_i(m)}{n} - \frac{\mathbb{E}T(m)}{m}\right| \geq \frac{\delta}{4}\right)\Bigg\} \bigcup \left\{\left|\frac{\mathbb{E}T(m)}{m} - \frac{\mathbb{E}T(n)}{n}\right| \geq \frac{\delta}{4}\right\}\Bigg), \tag{12}$$
$$\leq \lim_{n\to\infty} \mathbb{P}\left(\left|\frac{\sum_{i=1}^{\frac{n}{m}} T_i(m)}{n} - \frac{\mathbb{E}T(m)}{m}\right| \geq \frac{\delta}{4}\right) + \lim_{n\to\infty} \mathbb{P}\left(\left|\frac{\mathbb{E}T(m)}{m} - \frac{\mathbb{E}T(n)}{n}\right| \geq \frac{\delta}{4}\right) \tag{13}$$
$$\leq \lim_{n\to\infty} \mathbb{P}\left(\sum_{i=1}^{\frac{n}{m}} \left|T_i(m) - \mathbb{E}T(m)\right| \geq \frac{n\delta}{4}\right) + \lim_{n\to\infty} \mathbb{P}\left(\left|\frac{\mathbb{E}T(m)}{m} - \frac{\mathbb{E}T(n)}{n}\right| \geq \frac{\delta}{4}\right). \tag{14}$$

Again by Theorem 1, both $\frac{\mathbb{E}T(m)}{m}$ and $\frac{\mathbb{E}T(n)}{n}$ converge to the same constant R_d^* when n and m goes to infinity. Therefore, the second term vanishes. To help understand the proof, let's explain the inequalities above. The inequality between line (11) and line (12) can be derived by using the fact that

$$\left\{\left|\left(\frac{\sum_{i=1}^{\frac{n}{m}} T_i(m)}{n} - \frac{\mathbb{E}T(m)}{m}\right) + \left(\frac{\mathbb{E}T(m)}{m} - \frac{\mathbb{E}T(n)}{n}\right)\right| \geq \frac{\delta}{2}\right\} \subset \left\{\left|\frac{\sum_{i=1}^{\frac{n}{m}} T_i(m)}{n} - \frac{\mathbb{E}T(m)}{m}\right| \geq \frac{\delta}{4}\right\} \bigcup \left\{\left|\frac{\mathbb{E}T(m)}{m} - \frac{\mathbb{E}T(n)}{n}\right| \geq \frac{\delta}{4}\right\}$$

Union bound is applied between between line (12) and line (13). Line (13) and line (14) comes from the simple fact that

$$\left|\frac{\sum_{i=1}^{\frac{n}{m}} T_i(m)}{n} - \frac{\mathbb{E}T(m)}{m}\right| \leq \sum_{i=1}^{\frac{n}{m}} \left|T_i(m) - \mathbb{E}T(m)\right|.$$

Now let's use Lemma 4. Let $\mathcal{I}$ be the collection of nodes in the lattice. Define $\mathcal{C} = \{\mathcal{N}(\pi), \pi \in \Pi\}$, where $\mathcal{N}(\pi) = \{\mathbf{v} \in \pi\}$ is the set of nodes in the path π. Then, for the maximum-reward path with at most n steps, the maximum cardinality is $\max_{C\in\mathcal{C}} |C| \leq n$. Then, by substituting $T(n)$ for Z in Lemma 4,

$$P(|T(n) - \mathbb{E}[T(n)]| \geq u) \;\leq\; \exp\Big(-\frac{u^2}{64nL^2} + 64\Big).$$

Using the inequality we just derived and setting $m = c\log n$, we obtain

$$= \lim_{n\to\infty} \mathbb{P}\left(|R_{\text{iterative}}(n,m) - R_d^*| \geq \delta\right)$$
$$\leq \lim_{n\to\infty} \mathbb{P}\Big(\sum_{i=1}^{n/m} \big|\, T_i(m) - \mathbb{E}[T(m)]\,\big| \geq \frac{n\delta}{4}\Big) \tag{15}$$
$$\leq \lim_{n\to\infty} \sum_{i=1}^{\frac{n}{m}} \mathbb{P}\left(\big|\, T_i(m) - \mathbb{E}[T(m)]\,\big| \geq \frac{m\delta}{4}\right) \tag{16}$$
$$\leq \lim_{n\to\infty} \frac{n}{m} \cdot \exp\left(-\frac{(\frac{m\delta}{4})^2}{64mL^2} + 64\right) \tag{17}$$
$$= \lim_{n\to\infty} \frac{1}{c\log n} \cdot \exp\left(\left(1 - \frac{\delta^2}{1024L^2}\cdot c\right)\log n + 64\right) \tag{18}$$

The first inequality comes from line (14). Union bound is again applied between between line (15) and line (16). Lemma 4 is applied in line (17). Line(18) converges to 0 when the constant c is sufficiently large, *i.e.*, for any constant $c \geq \left(\frac{16L}{\delta}\right)^2$. ■

B. Proof for the Light-tailed Problem

We proceed to prove results for the original problem with light-tailed reward (*i.e.*, Theorem 4) by applying concentration of measure. Let's first introduce a lemma.

Lemma 5 (See [69], [89], [90])**.** *Assume Equation 5 holds true. Then there exists $c = c(d) < \infty$ such that for all $L > 0$ and all $\mathbf{x} < \mathbf{1}$, the shape function defined in Proposition 4 satisfies*

$$g(\mathbf{x}) - g^{(L)}(\mathbf{x}) \leq c\int_L^\infty (1 - F(x))^{1/d} dx.$$

Similar to the proof for Lemma 3, we will be using union bounds to break down the terms in the original inequality. More specifically,

$$\begin{aligned}
&\mathbb{P}\left(|R_{\text{iterative}}(n,m) - \mathbf{R}_d^*| \geq \delta\right)\\
=&\mathbb{P}\big(\big|R_{\text{iterative}}(n,m) - R^{(L)}_{\text{iterative}}(n,m) + R^{(L)}_{\text{iterative}}(n,m)\\
&\qquad - \mathbf{R}_d^* + \mathbf{R}_d^{*(L)} - \mathbf{R}_d^{*(L)}\big| \geq \delta\big)\\
\leq&\mathbb{P}\left(\left|R_{\text{iterative}}(n,m) - R^{(L)}_{\text{iterative}}(n,m)\right| \geq \delta/3\right)\\
&+ \mathbb{P}\left(\left|R^{(L)}_{\text{iterative}}(n,m) - \mathbf{R}_d^{*(L)}\right| \geq \delta/3\right)\\
&+ \mathbb{P}\left(\left|\mathbf{R}_d^* - \mathbf{R}_d^{*(L)}\right| \geq \delta/3\right)
\end{aligned}$$

The first and the third terms are the gap between the truncated and the original problems. By Lemma 5 both terms can be made arbitrarily small by choosing a large truncation threshold L. The second term also vanishes due to Lemma 3, when the sensing range is of order $O(\log n)$, as n increases. Therefore, with proper choice of the threshold L and sensing range $m = O(\log n)$,

$$\lim_{n\to\infty} \mathbb{P}\left(|R_{\text{iterative}}(n,m) - \mathbf{R}_d^*| \geq \delta\right) = 0.$$

APPENDIX F
PROOF FOR THEOREM 5

Recall that $T_{\text{iterative}}(n;m)$ is defined as $\sum_{i=1}^{n/m} T_i$ in Equation 6, where T_i is the reward collected at the i^{th} iteration in a receding horizon manner. We can apply the same argument in Appendix J to each T_i and derive that

$$T_i = O(m^{2/\alpha}).$$

Therefore, we obtain

$$R_{\text{iterative}}(n;m) = \frac{\sum_{i=1}^{n/m} T_i(m)}{n} = O\left(m^{2/\alpha-1}\right)$$

Since we already have $R_2^*(n) = O(n^{2/\alpha-1})$ from Proposition 2, it immediately follows that when $m = M(n)$ where $M(n)$ is a sub-linear function of n, we have

$$\lim_{n\to\infty} \frac{R_{\text{iterative}}(n;M(n))}{R_2^*(n)} = \lim_{n\to\infty} \left(\frac{M(n)}{n}\right)^{2/\alpha-1} = 0$$

APPENDIX G
PROOF FOR THEOREM 6

The proof applies a similar argument to Appendix A, using the lemma below (a continuous counterpart of Fekete's subadditive lemma).

Lemma 6 (See [91])**.** *For every measurable subadditive function $f : (0,\infty) \to \mathbb{R}$, the limit $\lim_{t\to\infty} \frac{f(t)}{t}$ exists and is equal to $\inf_{t>0} \frac{f(t)}{t}$.*

The function $\mathbb{E}[T_d^*(L)]$ is super-additive and by using the above lemma, it follows that the limit $\lim_{L\to\infty} \frac{\mathbb{E}[T_d^*(L)]}{L}$ exists and is equal to $\sup_L \frac{\mathbb{E}T_d^*(L)}{L}$.

APPENDIX H
PROOF FOR THEOREM 8

Let's denote the reachable region of the robot as $\mathcal{R}$. First note that

$$T_2^*(L) \geq \max_{\mathbf{v}\in\mathcal{R}} r(\mathbf{v}),$$

for all $\mathbf{v} \in \mathcal{R}$. If $\mathbb{E}r^d = \infty$, then for arbitrary $c > 0$ and infinitely many L, there must exist $\mathbf{v}$ such that $r(\mathbf{v}) \geq cL$. In other words, there exists infinitely many L such that for any c,

$$R_2^*(L) = \frac{T_2^*(L)}{L} \geq \frac{\max_{\mathbf{v}\in\mathcal{R}} r(\mathbf{v})}{L} \geq c$$

Therefore, $R_2^*(L)$ does not concentrate, and thus $\mathbf{R}_2^* = \infty$.

APPENDIX I
PROOF FOR PROPOSITION 3

To analyze the continuous problem, we approximate the continuous reward field with a two-dimensional $N \times N$ regular lattice. This approximation turns the continuous problem into a discrete problem that we are already familiar with from previous discussions.

More specifically, consider the set $\{\frac{1}{n} * L, \frac{2}{n} * L, \ldots, \frac{n-1}{n} * L, L\}^2 \subset [0,L]^2$. Let the node $\mathbf{v} = (i,j)$ on this discrete set be associated with $r_{i,j}^{(N)}$, which is the total reward that falls within the region $[\frac{i}{n} * L, \frac{i+1}{n} * L) \times [\frac{j}{n} * L, \frac{j+1}{n} * L)$ in the continuous model. This regions has an area of $s = \left(\frac{L}{N}\right)^2$.

Since the reward locations are distributed randomly in the field according to a Poisson process with parameter λ, the number of targets within any of such regions follows a Poisson distribution with intensity $p = \lambda \cdot s = \lambda\left(\frac{L}{N}\right)^2$. If we denote number of rewards within this region with a random variable $K \sim \mathbf{Pois}(p)$, then the total reward is

$$r_{i,j}^{(N)} = \sum_{k=1}^{K} r_k,$$

where r_k is drawn i.i.d. from the common reward distribution F.

Note that on this discrete set, all previous results (Theorems 1-5) apply. It remains to show that the discrete set is indeed a good approximate of the continuous model, and in fact this has been given in the proof for Theorem 2.1 in [66] for Pareto distributions. We rephrase the result in our notation as follows:

Theorem 14. *$T_2^*(1)$ is almost surely finite and as $n \to \infty$*

$$\frac{T_{(n)}}{n^{2/\alpha}} \to T_2^*(1) \quad \textit{in distribution},$$

where $T_{(n)}$ is the total reward on the approximate lattice.

Again, by the Skorohod Representation Theorem, we can define all the variables on the same probability space in such a way that the convergence occurs almost surely.

It is easy to see that for arbitrary size $L > 0$, the optimal reward on the continuous space can be approximated by $T_{(Ln)}$. Therefore by Proposition 3 we have

$$T_2^*(L) = O(L^{2/\alpha-1}).$$

APPENDIX J
PROOF FOR THEOREM 11

Proof: When $\alpha = 1$, the robot is operating on an isosceles right triangle region, as shown in Figure 11a. However, if $\alpha \neq 1$, the robot is operating on some isosceles triangle with the vertex angle being $2\arctan\alpha$, as shown in Figure 11b.

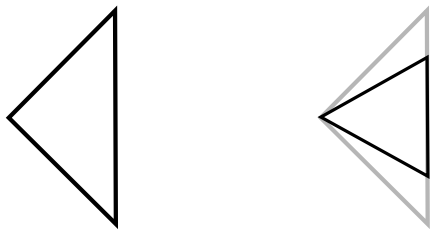

(a) agility $\alpha = 1$ (b) agility $\alpha \neq 1$

Fig. 11. The reachable set of the robot with drift, whose dynamics is described mathematically by Equations 1.

We extend all previous results to this isosceles triangle region by using the invariance property of Poisson point processes. By undergoing a deterministic transformation of the Poisson point process, we show the equivalence between this new problem (with non-unit α) and the one with $\alpha = 1$.

More specifically, this transformation is characterized by the mapping theorem of Poisson point process in [92]. Formally, let $f : \mathbb{R}^2 \to \mathbb{R}^2$ be the transformation between two poisson point processes. If f is an affine transformation

$$f(x) = Ax + b,$$

with $A \in \mathbb{R}^{2\times 2}$ being invertible, then the transformed Poisson Poisson has the new intensity

$$\nu(y) = \frac{1}{|A|}\lambda\left(f^{-1}(y)\right)\left(A^{-1}(y-b)\right),$$

where λ is the intensity of the original poisson point process.

In this case, the transformation from an isosceles triangle to an isosceles right triangle is linear. Therefore, $f = Ax$, where

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1/\alpha \end{bmatrix}.$$

Geometrically, this transformation stretches (or compresses) the isosceles triangle vertically until the vertex angle is 90 degrees, as shown in Figure 11b. Note that we assume λ is constant in this problem, and therefore the new intensity is $\nu = \alpha \cdot \lambda$. Therefore, the maximum-reward motion problem in $\mathbb{R}^2$, where a robot has agility α and the reward is generated by a Poisson point process with intensity λ, is equivalent to the maximum-reward motion problem with agility 1 and intensity $\alpha\lambda$, with respect to the mean reward collected.

Additionally, since the mean reward $\mathbb{E}[T(L)]$ grows proportionally with $\sqrt{\lambda}$, we conclude that it is also linear with $\sqrt{\alpha}$, *i.e.*,

$$\mathbb{E}[T(L)] = c\sqrt{\alpha} = c\sqrt{w/v}.$$

for some constant $c > 0$. ■